

BEYOND VALUE ITERATION FOR PARITY GAMES: STRATEGY ITERATION WITH UNIVERSAL TREES*

ZHUAN KHYE KOH ^a AND GEORG LOHO ^b

^a Boston University, 665 Commonwealth Avenue, Boston, MA 02215, USA.
e-mail address: zkkoh@bu.edu

^b University of Twente, Drienerlolaan 5, 7522 NB Enschede, The Netherlands.
e-mail address: g.loho@utwente.nl

ABSTRACT. Parity games have witnessed several new quasi-polynomial algorithms since the breakthrough result of Calude et al. (STOC 2017). The combinatorial object underlying these approaches is a *universal tree*, as identified by Czerwiński et al. (SODA 2019). By proving a quasi-polynomial lower bound on the size of a universal tree, they have highlighted a barrier that must be overcome by all existing approaches to attain polynomial running time. This is due to the existence of worst case instances which force these algorithms to explore a large portion of the tree.

As an attempt to overcome this barrier, we propose a strategy iteration framework which can be applied on any universal tree. It is at least as fast as its value iteration counterparts, while allowing one to take bigger leaps in the universal tree. Our main technical contribution is an efficient method for computing the least fixed point of 1-player games. This is achieved via a careful adaptation of shortest path algorithms to the setting of ordered trees. By plugging in the universal tree of Jurdziński and Lazić (LICS 2017), or the Strahler universal tree of Daviaud et al. (ICALP 2020), we obtain instantiations of the general framework that take time $O(mn^2 \log n \log d)$ and $O(mn^2 \log^3 n \log d)$ respectively per iteration.

1. INTRODUCTION

A *parity game* is an infinite duration game between two players Even and Odd. It takes place on a sinkless directed graph $G = (V, E)$ equipped with a *priority* function $\pi : V \rightarrow \{1, 2, \dots, d\}$. Let $n = |V|$ and $m = |E|$. The node set V is partitioned into $V_0 \sqcup V_1$ such that nodes in V_0 and V_1 are owned by Even and Odd respectively. The game starts when a token is placed on a node. In each turn, the owner of the current node moves the token along an

Key words and phrases: parity games, strategy iteration, value iteration, progress measure, universal trees.

* An extended abstract of this paper has appeared in Proceedings of the 47th International Symposium on Mathematical Foundations of Computer Science, MFCS 2022.

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement nos. 757481–ScaleOpt and 805241–QIP). Z. K. Koh—This work was done while the author was at the London School of Economics and Centrum Wiskunde & Informatica.

outgoing arc to the next node, resulting in an infinite walk. If the highest priority occurring infinitely often in this walk is even, then Even wins. Otherwise, Odd wins.

By the positional determinacy of parity games [Mar75, EJ91], there exists a partition of V into two subsets from which Even and Odd can force a win respectively. The main algorithmic problem of parity games is to determine this partition, or equivalently, to decide the winner given a starting node. This is a notorious problem that lies in $\text{NP} \cap \text{co-NP}$ [EJS93], and also in $\text{UP} \cap \text{co-UP}$ [Jur98], with no known polynomial algorithm to date.

Due to its intriguing complexity status, as well as its fundamental role in automata theory and logic [EJS93, KV98], parity games have been intensely studied over the past three decades. Prior to 2017, algorithms for solving parity games, e.g. [Zie98, Jur00, VJ00, BSV03, Sch08, JPZ08, Sch17, MRR18, BDM18], are either exponential or mildly subexponential. In a breakthrough result, Calude et al. [CJK⁺17] gave the first quasi-polynomial algorithm. Since then, many other quasi-polynomial algorithms [FJdK⁺19, JL17, Leh18, Par19, BDM⁺25] have been developed. Most of them have been unified by Czerwiński et al. [CDF⁺19] via the concept of a *universal tree*. A universal tree is an ordered tree into which every ordered tree of a certain size can be isomorphically embedded. They proved a quasi-polynomial lower bound on the size of a universal tree.

Value iteration. The starting point of this paper is the classic *progress measure* algorithm [Jur00, JL17] for solving parity games. It belongs to a broad class of algorithms called *value iteration* – a well-known method for solving more general games on graphs such as mean payoff games and stochastic games. In value iteration, every node v in G is assigned a value $\mu(v) \in \mathcal{V}$ from some totally ordered set \mathcal{V} , and the values are locally improved until we reach the *least fixed point* of a set of operators associated with the game. The set \mathcal{V} is called the *value domain*, which is usually a bounded set of real numbers or integers. For the progress measure algorithm, its value domain is the set of leaves $L(T)$ in a universal tree T . As the values are monotonically improved, the running time is proportional to $|L(T)|$. The first progress measure algorithm of Jurdziński [Jur00] uses a perfect n -ary tree, which runs in exponential time. Its subsequent improvement by Jurdziński and Lazić [JL17] uses a quasi-polynomial-sized tree, which runs in $n^{\log(d/\log n)+O(1)}$ time.

Despite having good theoretical efficiency, the progress measure algorithm is not robust against its worst-case behaviour. In fact, it is known to realize its worst-case running time on very simple instances. As an example, let (G, π) be an arbitrary instance with maximum priority d , with d being even. For a small odd constant k , if we add two nodes of priority k as shown in Figure 1, then the progress measure algorithm realizes its worst-case running time. This is because the values of these nodes are updated superpolynomially many times.



Figure 1: A worst-case construction for the progress measure algorithm. Nodes in V_0 and V_1 are drawn as squares and circles, respectively.

Strategy iteration. A different but related method for solving games on graphs is *strategy iteration*. For a parity game (G, π) , a (*positional*) *strategy* τ for a player (say Odd) is a choice of an outgoing arc from every node in V_1 . Removing the unchosen outgoing arcs from every node in V_1 results in a *strategy subgraph* $G_\tau \subseteq G$. A general framework for strategy iteration

is given, e.g., in [Fri11]. Following that exposition, one fixes a suitable value domain \mathcal{V} and associates a valuation $\mu : V \rightarrow \mathcal{V}$ to each strategy for Odd. The valuation of a strategy τ can be thought of as an evaluation of the best counterstrategy for Even. More formally, it is a fixed point of a set of operators associated with the 1-player game (G_τ, π) for Even. Given two strategies τ, τ' and their respective valuations μ, μ' , we say that τ' is *better* than τ if $\mu'(v) \geq \mu(v)$ for all $v \in V$ and $\mu' \neq \mu$. In every iteration, strategy iteration maintains a strategy τ for Odd and its corresponding valuation μ . Based on μ and a *pivot rule*, it switches to a better strategy τ' , and then computes its valuation μ' . This process is repeated until we reach an optimal strategy for Odd.

Originally introduced by Hoffman and Karp for stochastic games [HK66], variants of strategy iteration for parity games have been developed [Pur95, VJ00, BSV03, BV07, Sch08, Lut08]. They usually perform well in practice, but tedious constructions of their worst case (sub)exponential complexity are known [Fri09]. Motivated by the construction of quasi-polynomial universal trees [JL17, DJT20], a natural question is whether there exists a strategy iteration algorithm with value domain $L(T)$ for a universal tree T . It is not hard to see that with value domain $L(T)$, unfortunately, the fixed point of a 1-player game (G_τ, π) may not be unique. Moreover, in a recent thesis [Ohl21], Ohlmann showed that a valuation that is fit for strategy iteration cannot be defined using $L(T)$ for the aforementioned trees [Jur00, JL17, DJT20].

Our contribution. We show that an adaptation of strategy iteration with value domain $L(T)$ is still possible. To circumvent the impossibility result of Ohlmann [Ohl21], we slightly alter the strategy iteration framework as follows. After pivoting to a strategy τ' in an iteration, we update the current node labeling μ to the least fixed point of $(G_{\tau'}, \pi)$ that is *pointwise at least* μ . In other words, we force μ to increase (whereas this happens automatically in the previous framework). Since the fixed point of a 1-player game may not be unique, this means that we may encounter a strategy more than once during the course of the algorithm. The motivation of our approach comes from tropical geometry, as discussed in Appendix A.

To carry out each iteration efficiently, we give a combinatorial method for computing the least fixed point of 1-player games with value domain $L(T)$. In the literature, valuations are usually computed using shortest path techniques, such as label-setting (e.g. Dijkstra's algorithm) and label-correcting (e.g. Bellman–Ford algorithm), with running times $O(m + n \log n)$ and $O(mn)$ respectively. We adapt these techniques to the setting of ordered trees. When T is instantiated as a specific universal tree constructed in the literature, we obtain the following running times:

- The perfect n -ary tree of height $d/2$ takes $O(d(m + n \log n))$.
- The universal tree of Jurdziński and Lazić [JL17] takes $O(mn^2 \log n \log d)$.
- The Strahler universal tree of Daviaud et al. [DJT20] takes $O(mn^2 \log^3 n \log d)$.

The total number of strategy iterations is trivially bounded by $n|L(T)|$, the same bound for the progress measure algorithm. Whereas we do not obtain a strict improvement over previous running time bounds, it is conceivable that our algorithm would terminate in fewer iterations than the progress measure algorithm on most examples. Moreover, our framework provides large flexibility in the choice of pivot rules. Identifying a pivot rule that may provide strictly improved (and possibly even polynomial) running time is left for future research.

1.1. Computing the Least Fixed Point of 1-Player Games. Let (G_τ, π) be a 1-player game for Even, and μ^* be its least fixed point with value domain $L(T)$ for some universal tree T . Starting from $\mu(v) = \min L(T)$ for all $v \in V$, the progress measure algorithm successively lifts the label of a node based on the labels of its out-neighbours until μ^* is reached. This procedure is not polynomial in general, even on 1-player games. So, instead of approaching μ^* from below, we approach it from above. This is reminiscent of shortest path algorithms, where node labels form upper bounds on the shortest path distances throughout the algorithm. To compute shortest paths to a target node t , the label at t is initialized to 0, while the labels at all other nodes are initialized to $+\infty$. In a label-correcting method like the Bellman–Ford algorithm, the labels are monotonically decreased until they converge to shortest path distances. Differently, in a label-setting method such as Dijkstra’s algorithm, the label of a node is directly fixed to its shortest path distance in every iteration. We refer to Ahuja et al. [AMO93] for an overview on label-correcting and label-setting techniques for the shortest path problem.

There are two main challenges when applying these techniques to the setting of ordered trees. First, the target node t is not explicitly given. Let us call a cycle *even* if its maximum priority is even. By inspecting the operators associated with (G_τ, π) , one can deduce that for every node v , one gets $\mu^*(v) < \top$ only if v can reach an even cycle in G_τ (Lemma 2.4). However, it is unclear which even cycle determines the value of $\mu^*(v)$ a priori. To overcome this issue, we consider maximum priority nodes in even cycles of G_τ , which we call *base nodes*. They will play the role of target nodes.

The second challenge is how to initialize the labels on the base nodes. When T is a perfect n -ary tree, one can show that for every base node w , the label $\mu^*(w)$ is equal to the smallest leaf in T (Theorem 4.5 and Remark 4.6). This property of perfect n -ary trees allows us to compute μ^* using a label-setting method, with the initial node labeling ν given by $\nu(v) = \min L(T)$ if v is a base node, and $\nu(v) = \top$ otherwise. Here, \top is defined to be bigger than every element in $L(T)$ (so \top is analogous to $+\infty$ for real numbers).

For the shortest path problem, Dijkstra’s algorithm selects a node with the smallest label to be fixed in every iteration. Unfortunately, this criterion does not work for us because the representation of a parity game as a mean payoff game can have negative arc weights, and so the correctness of Dijkstra’s algorithm is not guaranteed. To fix this issue, we define a potential function to guide the algorithm in selecting the correct node. Let H be the subgraph of G_τ obtained by deleting all the base nodes. For $p \in \mathbb{N}$, let H_p be the subgraph of H induced by the nodes with priority at most p . The potential function is constructed by interlacing the label on a node with a tuple that encodes its topological orders in H_2, H_4, \dots . In every iteration, a node with the smallest potential is selected, and its label is fixed.

When T is a general universal tree, we lose the property that $\mu^*(w) = \min L(T)$ for all base nodes w . Since we do not know $\mu^*(w)$, we resort to label-correcting techniques. It turns out that the Bellman–Ford algorithm has a less stringent requirement on the initial node labeling ν . For every base node w , we identify a range $[\mu^*(w), \hat{\mu}(w)]$ for $\nu(w)$ such that the Bellman–Ford algorithm is guaranteed to return μ^* .

To understand the upper bound $\hat{\mu}(w)$, consider the set of cycles in G_τ which contain w as a maximum priority node. Clearly, every such cycle C induces a subgame (C, π) on which Even wins because $\pi(w)$ is even. The least fixed point of (C, π) consists of leaves in T which belong to a subtree T_C of height $j := \pi(w)/2$. The upper bound $\hat{\mu}(w)$ is given by a cycle C with the ‘narrowest’ T_C . Since the subtrees of height j in T form a poset $(\mathcal{T}_j, \sqsubseteq)$

with respect to the partial order of embeddability, the notion of ‘width’ here is defined using a chain cover \mathcal{C}_j of $(\mathcal{T}_j, \sqsubseteq)$.

To compute $\nu(w) \in [\mu^*(w), \widehat{\mu}(w)]$ for a base node w , we construct an arc-weighted auxiliary digraph D on the set of base nodes. Every arc uv in D represents a path from base node u to base node v in G_τ . We show that a minimum bottleneck cycle containing w in D gives rise to a cycle C in G_τ , such that the smallest leaf in T_C lies in $[\mu^*(w), \widehat{\mu}(w)]$. After computing $\nu(w)$ for all base nodes w , we set $\nu(v) = \top$ for the other nodes.

With this initial node labeling ν , the Bellman–Ford algorithm returns μ^* in $O(mn)$ steps. The overall running time is dominated by the computation of ν , whose running time is proportional to the size of the chain cover \mathcal{C}_j . We prove that the quasi-polynomial universal trees constructed in the literature [JL17, DJT20] admit small chain covers. Using this result, we then give efficient implementations of our method for these trees.

Paper organization. In Section 2, we introduce notation and provide the necessary preliminaries on parity games and universal trees. Section 3 contains our strategy iteration framework based on universal trees. Section 4 gives a label-setting method for computing the least fixed point of 1-player games, and demonstrates its applicability to perfect n -ary trees. In Section 5, we develop a label-correcting method for the same task, and apply it to known constructions of quasi-polynomial universal trees.

2. PRELIMINARIES ON PARITY GAMES AND UNIVERSAL TREES

For $d \in \mathbb{N}$, let $[d] = \{1, 2, \dots, d\}$. For a graph G , we use $V(G)$ as its vertex set, and $E(G)$ as its edge set. A parity game instance is given by (G, π) , where $G = (V, E)$ is a sinkless directed graph with $V = V_0 \sqcup V_1$, and $\pi : V \rightarrow [d]$ is a priority function. Recall that the players are Even and Odd, who control vertices in V_0 and V_1 respectively. Without loss of generality, we may assume that d is even. In this paper, we are only concerned with positional strategies. A *strategy* for Odd is a function $\tau : V_1 \rightarrow V$ such that $v\tau(v) \in E$ for all $v \in V_1$. Its *strategy subgraph* is $G_\tau = (V, E_\tau)$, where $E_\tau := \{vw \in E : v \in V_0\} \cup \{v\tau(v) : v \in V_1\}$. A strategy for Even and its strategy subgraph are defined analogously. We always denote a strategy for Even as σ , and a strategy for Odd as τ . If we fix a strategy τ for Odd, the resulting instance (G_τ, π) is a *1-player game* for Even.

For the sake of brevity, we overload the priority function π as follows. Given a subgraph $H \subseteq G$, let $\pi(H)$ be the highest priority in H . The subgraph H is said to be *even* if $\pi(H)$ is even, and *odd* otherwise. For a fixed π , we denote by $\Pi(H)$ the set of nodes with the highest priority in H . If $v \in \Pi(H)$, we say that v *dominates* H . For $p \in [d]$, H_p refers to the subgraph of H induced by nodes with priority at most p . For a node v , let $\delta_H^-(v)$ and $\delta_H^+(v)$ be the incoming and outgoing arcs of v in H respectively. Similarly, let $N_H^-(v)$ and $N_H^+(v)$ be the in-neighbors and out-neighbors of v in H respectively. When H is clear from context, we will omit it from the subscripts.

The win of a player can be certified by *node labels* from a *universal tree*, as stated in Theorem 2.3. We give the necessary background for this now.

2.1. Ordered Trees and Universal Trees. An *ordered tree* T is a prefix-closed set of tuples, whose elements are drawn from a linearly ordered set M . The linear order of M lexicographically extends to T . We note that any proper prefix of a tuple is smaller than the tuple itself. Equivalently, T can be thought of as a rooted tree, whose root we denote by

r . Under this interpretation, elements in M correspond to the branching directions at each vertex of T (see Figures 2 and 3 for examples). Every tuple then corresponds to a vertex $v \in V(T)$. This is because the tuple can be read by traversing the unique r - v path in T . Observe that v is an h -tuple if and only if v is at depth h in T . In particular, r is the empty tuple.

In this paper, we always use the terms ‘vertex’ and ‘edge’ when referring to an ordered tree T . The terms ‘node’ and ‘arc’ are reserved for the game graph G .

Given an ordered tree T of height h , let $L(T)$ be the set of leaves in T . For convenience, we assume that every leaf in T is at depth h throughout. The tuple representing a leaf $\xi \in L(T)$ is denoted as $\xi = (\xi_{2h-1}, \xi_{2h-3}, \dots, \xi_1)$, where $\xi_i \in M$ for all i . We refer to ξ_{2h-1} as the *first* component of ξ , even though it has index $2h - 1$. For a fixed $p \in [2h]$, the p -truncation of ξ is

$$\xi|_p := \begin{cases} (\xi_{2h-1}, \xi_{2h-3}, \dots, \xi_{p+1}), & \text{if } p \text{ is even} \\ (\xi_{2h-1}, \xi_{2h-3}, \dots, \xi_p), & \text{if } p \text{ is odd.} \end{cases}$$

In other words, the p -truncation of a tuple is obtained by deleting the components with index less than p . Note that a truncated tuple is an ancestor of the untruncated tuple in T .

Definition 2.1. Given ordered trees T and T' , we say that T *embeds into* T' (denoted $T \sqsubseteq T'$) if there exists an *injective* and *order-preserving* homomorphism from T to T' such that leaves in T are mapped to leaves in T' . Formally, this is an injective function $f : V(T) \rightarrow V(T')$ which satisfies the following properties:

- (1) For all $u, v \in V(T)$, $uv \in E(T)$ implies $f(u)f(v) \in E(T')$;
- (2) For all $u, v \in V(T)$, $u \leq v$ implies $f(u) \leq f(v)$.
- (3) $f(u) \in L(T')$ for all $u \in L(T)$.

We write $T \equiv T'$ if $T \sqsubseteq T'$ and $T' \sqsubseteq T$. Also, $T \sqsubset T'$ if $T \sqsubseteq T'$ and $T \not\equiv T'$.

In the definition above, since f is order-preserving, the children of every vertex in T are mapped to the children of its image injectively such that their order is preserved. As an example, the tree in Figure 3 embeds into the tree in Figure 2. It is easy to verify that \sqsubseteq is a partial order on the set of all ordered trees.

Definition 2.2. An (ℓ, h) -*universal tree* is an ordered tree T' of height h such that $T \sqsubseteq T'$ for every ordered tree T of height h and with at most ℓ leaves, all at depth exactly h .

The simplest example of an (ℓ, h) -universal tree is the perfect ℓ -ary tree of height h , which we call a *perfect universal tree*. The linearly ordered set M for this tree can be chosen as $\{0, 1, \dots, \ell - 1\}$ (see Figure 2 for an example). It has ℓ^h leaves, which grows exponentially with h . Jurdziński and Lazić [JL17] constructed an (ℓ, h) -universal tree with at most $\ell^{\log h + O(1)}$ leaves, which we call a *succinct universal tree*. In this tree, every leaf ξ corresponds to an h -tuple of binary strings with at most $\lceil \log(\ell) \rceil$ bits in total¹. We use $|\xi|$ and $|\xi_i|$ to denote the total number of bits in ξ and ξ_i respectively. The linearly ordered set M for this tree consists of finite binary strings, where $\varepsilon \in M$ is the empty string (see Figure 3 for an example). For any pair of binary strings $s, s' \in M$ and a bit b , the linear order on M is defined as $0s < \varepsilon < 1s'$ and $bs < bs' \iff s < s'$.

¹A slightly looser bound of $\lceil \log \ell \rceil$ was derived in [JL17, Lemma 1]. It can be strengthened to $\lfloor \log \ell \rfloor$ with virtually no change in the proof.




Figure 2: The perfect (3,2)-universal tree. Figure 3: The succinct (3,2)-universal tree.

2.2. Node Labelings from Universal Trees. Let (G, π) be a parity game instance and T be an ordered tree of height $d/2$. We augment the set of leaves with an extra *top* element \top , denoted $\bar{L}(T) := L(T) \cup \{\top\}$, such that $\top > v$ for all $v \in V(T)$. In this way, \top is analogous to $+\infty$ for real numbers. We also set $\top|_p := \top$ for all $p \in [d]$. A function $\mu : V \rightarrow \bar{L}(T)$ which maps the nodes in G to $\bar{L}(T)$ is called a *node labeling*. The node labeling is *finite* if $\mu(v) \neq \top$ for all $v \in V$.

For a subgraph H of G , we say that μ is *feasible in H* if there exists a strategy $\sigma : V_0 \rightarrow V$ for Even with $v\sigma(v) \in E(H)$ whenever $\delta_H^+(v) \neq \emptyset$, such that the following condition holds for every arc vw in $H \cap G_\sigma$:

- If $\pi(v)$ is even, then $\mu(v)|_{\pi(v)} \geq \mu(w)|_{\pi(v)}$.
- If $\pi(v)$ is odd, then $\mu(v)|_{\pi(v)} > \mu(w)|_{\pi(v)}$ or $\mu(v) = \mu(w) = \top$.

An arc vw which does not satisfy the condition above is called *violated* (with respect to μ). On the other hand, we distinguish two possibilities for non-violated arcs. If $\mu(v)$ is the smallest element in $\bar{L}(T)$ such that vw is not violated, then vw is said to be *tight*. Any arc which is neither tight nor violated is called *loose*. We remark that loose does not mean non-tight. We say that a subgraph is *tight* if it consists of tight arcs.

In the literature, a node labeling which is feasible in G is also called a *progress measure*. The node labeling given by $\mu(v) = \top$ for all $v \in V$ is trivially feasible in G . However, we are primarily interested in progress measures with maximal finite support, i.e. the set of nodes with non-top labels is inclusion-wise maximal.

Theorem 2.3 [Jur00]. *Given an $(n, d/2)$ -universal tree T , let $\mu^* : V \rightarrow \bar{L}(T)$ be a node labeling which is feasible in G and has maximal finite support. Then, Even wins from $v \in V$ if and only if $\mu^*(v) \neq \top$.*

The above theorem formalizes the following intuition: nodes with smaller labels are more advantageous for Even to play on. Note that if μ is a minimal node labeling which is feasible in G , i.e. μ' is infeasible in G for all $\mu' < \mu$, then there exists a strategy σ for Even such that $v\sigma(v)$ is tight for all $v \in V_0$. The next observation is well-known (see, e.g., [JL17, Lemma 2]) and follows directly from the definition of feasibility.

Lemma 2.4 (Cycle Lemma). *If a node labeling μ is feasible in a cycle C , then*

$$\mu(v)|_{\pi(C)} = \mu(w)|_{\pi(C)} \text{ for all } v, w \in V(C) .$$

Furthermore, if $\mu(v) \neq \top$ for some $v \in V(C)$, then C is even.

We assume to have access to the following algorithmic primitive TIGHTEN, whose running time we denote by $\gamma(T)$. Its implementation depends on the ordered tree T . For instance, $\gamma(T) = O(d)$ if T is a perfect $(n, d/2)$ -universal tree. If T is a succinct $(n, d/2)$ -universal tree, Jurdziński and Lazić [JL17, Theorem 7] showed that $\gamma(T) = O(\log n \log d)$.

$\text{TIGHTEN}(\mu, vw)$

Given a node labeling $\mu : V \rightarrow \bar{L}(T)$ and an arc $vw \in E$, return the unique element $\xi \in \bar{L}(T)$ such that vw is tight after setting $\mu(v)$, the label of v , to ξ .

Given a node labeling $\mu : V \rightarrow \bar{L}(T)$ and an arc $vw \in E$, let $\text{lift}(\mu, vw)$ be the smallest element $\xi \in \bar{L}(T)$ such that $\xi \geq \mu(v)$ and vw is not violated after setting $\mu(v)$ to ξ . Observe that if vw is violated, $\text{lift}(\mu, vw)$ is given by $\text{TIGHTEN}(\mu, vw)$. Otherwise, it is equal to $\mu(v)$. Hence, it can be computed in $\gamma(T)$ time.

2.3. Fixed Points in Lattices. We recall a fundamental result on the existence of fixed points in lattices. Let \mathcal{L} be a non-empty finite lattice, and let $\mu, \nu \in \mathcal{L}$. An operator $\phi : \mathcal{L} \rightarrow \mathcal{L}$ is *monotone* if $\mu \leq \nu \Rightarrow \phi(\mu) \leq \phi(\nu)$, and it is *inflationary* if $\mu \leq \phi(\mu)$. Given a family \mathcal{G} of inflationary monotone operators on \mathcal{L} , we denote $\mu^{\mathcal{G}}$ as the *least* (simultaneous) fixed point of \mathcal{G} which is *at least* μ . The following proposition, which can be seen a variant of the Knaster–Tarski theorem, guarantees the existence of $\mu^{\mathcal{G}}$. We remark that $\mu^{\mathcal{G}}$ may not exist if the inflationary assumption is removed.

Proposition 2.5. *Let \mathcal{G} be a family of inflationary monotone operators on a non-empty finite lattice \mathcal{L} . For any $\mu \in \mathcal{L}$ and $\mathcal{H} \subseteq \mathcal{G}$,*

- (i) *The least fixed point $\mu^{\mathcal{H}}$ exists. It can be obtained by applying the operators in \mathcal{H} to μ in any order until convergence.*
- (ii) *The least fixed point is non-decreasing with the set of operators: $\mu^{\mathcal{H}} \leq \mu^{\mathcal{G}}$.*
- (iii) *The least fixed point is monotone: if $\mu \leq \nu$ then $\mu^{\mathcal{H}} \leq \nu^{\mathcal{H}}$.*

Proof. For part (i), let $\mu =: \mu_1 < \mu_2 < \dots < \mu_k$ be a maximal sequence obtained by applying the operators in \mathcal{H} to μ ; this sequence is finite because \mathcal{L} is finite. By maximality, μ_k is a simultaneous fixed point of \mathcal{H} . To show minimality of this fixed point μ_k , let $\mu' \geq \mu$ be an arbitrary simultaneous fixed point of \mathcal{H} . For each $i < k$, let $\phi_i \in \mathcal{H}$ where $\phi_i(\mu_i) = \mu_{i+1}$. Then, by monotonicity and induction, $\mu' = (\phi_i \circ \dots \circ \phi_1)(\mu') \geq \phi_i(\mu_i) = \mu_{i+1}$ for all $i < k$. In particular, $\mu' \geq \mu_k$, so $\mu_k = \mu^{\mathcal{H}}$.

Part (ii) holds by definition because any simultaneous fixed point of \mathcal{G} is a simultaneous fixed point of \mathcal{H} . Part (iii) holds again by definition because $\nu^{\mathcal{H}}$ is a simultaneous fixed point of \mathcal{H} which is at least μ . \square

An operator $\psi : \mathcal{L} \rightarrow \mathcal{L}$ is *deflationary* if $\mu \geq \psi(\mu)$. Considering the lattice through an order-reversing poset isomorphism implies that the analogous statements of Proposition 2.5 hold for deflationary monotone operators. Given a family \mathcal{G} of deflationary monotone operators on \mathcal{L} , we denote $\mu^{\mathcal{G}}$ as *greatest* (simultaneous) fixed point of \mathcal{G} which is *at most* μ .

In this paper, \mathcal{L} is the finite lattice of node labelings mapping V to $\bar{L}(T)$ for a universal tree T . The partial order of \mathcal{L} is induced by the total order of $\bar{L}(T)$. For a subgraph $H \subseteq G$, consider the following operators. For every node $v \in V_0$, define $\text{Lift}_v : \mathcal{L} \times V \rightarrow \bar{L}(T)$ as

$$\text{Lift}_v(\mu, u) := \begin{cases} \min_{vw \in E(H)} \text{lift}(\mu, vw), & \text{if } u = v \text{ and } \delta_H^+(v) \neq \emptyset \\ \mu(u), & \text{otherwise.} \end{cases}$$

For every arc $vw \in E(H)$ where $v \in V_1$, define $\text{Lift}_{vw} : \mathcal{L} \times V \rightarrow \bar{L}(T)$ as

$$\text{Lift}_{vw}(\mu, u) := \begin{cases} \text{lift}(\mu, vw), & \text{if } u = v \\ \mu(u), & \text{otherwise.} \end{cases}$$

We denote $\mathcal{H}^\uparrow := \{\text{Lift}_v : v \in V_0\} \cup \{\text{Lift}_{vw} : v \in V_1\}$ as the operators in H . Since they are inflationary and monotone, for any $\mu \in \mathcal{L}$, the least fixed point $\mu^{\mathcal{H}^\uparrow}$ exists by Proposition 2.5(i). Note that a node labeling is a fixed point of \mathcal{H}^\uparrow if and only if it is feasible in H . In this terminology, the *progress measure algorithm* [Jur00, JL17] is an iterative application of the operators in \mathcal{G}^\uparrow to μ to obtain $\mu^{\mathcal{G}^\uparrow}$.

3. STRATEGY ITERATION WITH TREE LABELS

In this section, we present a strategy iteration algorithm (Algorithm 1) whose pivots are guided by a universal tree. It takes as input an instance (G, π) , a universal tree T , and an initial strategy τ_1 for Odd. Throughout, it maintains a node labeling $\mu : V \rightarrow \bar{L}(T)$, initialized as the least simultaneous fixed point of $\mathcal{G}_{\tau_1}^\uparrow$. At the start of every iteration, the algorithm maintains a strategy τ for Odd such that the following two invariants hold: (1) μ is feasible in G_τ ; and (2) there are no loose arcs in G_τ with respect to μ . So, every arc in G_τ is either tight (usable by Even in her counterstrategy σ) or violated (not used by Even). Note that our initial node labeling satisfies these conditions with respect to τ_1 .

Algorithm 1 Strategy iteration with tree labels: (G, π) instance, T universal tree, τ_1 initial strategy for Odd

```

1: procedure STRATEGYITERATION( $(G, \pi), T, \tau_1$ )
2:    $\mu(v) \leftarrow \min L(T) \forall v \in V$ 
3:    $\tau \leftarrow \tau_1, \mu \leftarrow \mu^{\mathcal{G}_\tau^\uparrow}$ 
4:   while  $\exists$  an admissible arc in  $G$  with respect to  $\mu$  do
5:     Pivot to a strategy  $\tau'$  by selecting admissible arc(s)       $\triangleright$  requires a pivot rule
6:      $\tau \leftarrow \tau', \mu \leftarrow \mu^{\mathcal{G}_{\tau'}^\uparrow}$ 
7:   return  $\tau, \mu$ 

```

For $v \in V_1$, we call a violated arc $vw \in E$ with respect to μ *admissible* (as it admits Odd to perform an improvement). If there are no admissible arcs in G , then the algorithm terminates. In this case, μ is feasible in G . Otherwise, Odd pivots to a new strategy τ' by switching to admissible arc(s). The choice of which admissible arc(s) to pick is governed by a *pivot rule*. Then, μ is updated to $\mu^{\mathcal{G}_{\tau'}^\uparrow}$. Note that both invariants continue to hold at the start of the next iteration. Indeed, there are no loose arc in G_τ with respect to μ due to the minimality of $\mu^{\mathcal{G}_{\tau'}^\uparrow}$.

We remark that a strategy τ may occur more than once during the course of the algorithm, as mentioned in the description of strategy iteration in Section 1. This is because the fixed points of $\mathcal{G}_\tau^\uparrow$ are not necessarily unique. See Figure 4 for an example run with a perfect universal tree and a succinct universal tree.

The correctness of Algorithm 1 is an easy consequence of Proposition 2.5.

Theorem 3.1. *Algorithm 1 returns the pointwise minimal node labeling $\mu^* : V \rightarrow \bar{L}(T)$ which is feasible in G .*

Proof. The node labeling μ is monotone increasing in every iteration. Since \mathcal{L} is finite and the all-top node labeling is feasible in G , the algorithm terminates. Let μ^* be the pointwise minimal node labeling which is feasible in G . Note that μ^* is the least simultaneous fixed

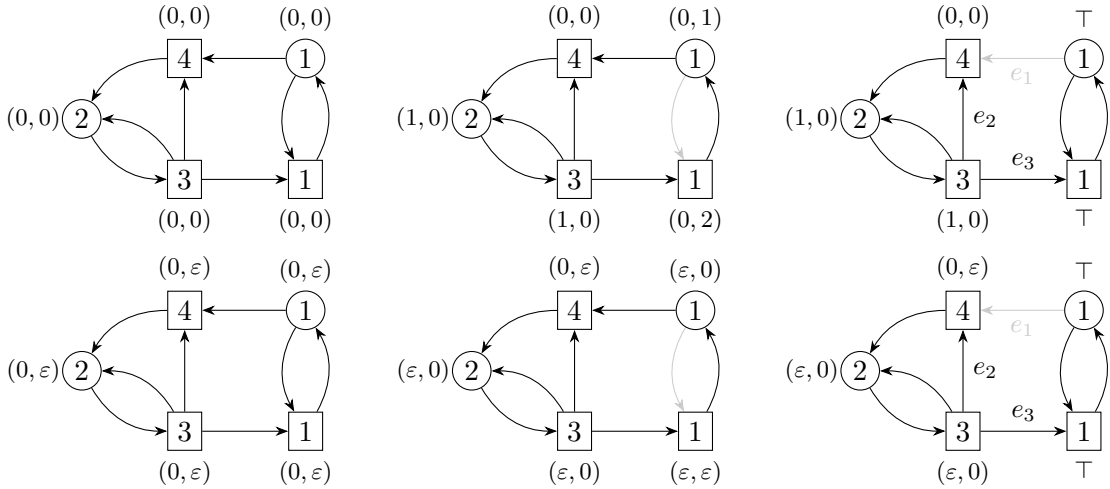


Figure 4: The top and bottom rows illustrate an example run of Algorithm 1 with the perfect (3,2)-universal tree and the succinct (3,2)-universal tree respectively. In each row, the left figure depicts the game instance (nodes in V_0 and V_1 are drawn as squares and circles respectively). The next two figures show Odd's strategy and the node labeling at the start of Iteration 1 and 2. The arcs not selected by Odd are greyed out. In the right figure, e_1 is loose, e_2 is tight, and e_3 is violated.

point of \mathcal{G}^\dagger in \mathcal{L} . By induction and Proposition 2.5 (ii), we have $\mu^{\mathcal{G}_\tau^\dagger} \leq \mu^*$ in every iteration. As the algorithm terminates with a simultaneous fixed point of \mathcal{G}^\dagger , it terminates with μ^* . \square

Thus, by Theorem 2.3, the algorithm correctly determines the winning positions for Even.

Running time. In Algorithm 1, a naive method for computing the least fixed point $\mu^{\mathcal{G}_\tau^\dagger}$ is to iterate the operators in \mathcal{G}_τ^\dagger on μ until convergence. The operators Lift_v and Lift_{vw} can be implemented to run in $O(|\delta^+(v)|\gamma(T))$ time. Since μ is monotone increasing throughout, the total running time of Algorithm 1 is

$$O\left(\sum_{v \in V} |\delta^+(v)|\gamma(T)|L(T)|\right) = O(m\gamma(T)|L(T)|),$$

which matches progress measure algorithms [Jur00, JL17]. However, this method of computing $\mu^{\mathcal{G}_\tau^\dagger}$ can take $\Omega(\gamma(T)|L(T)|)$ time; recall that $|L(T)|$ is at least quasi-polynomial for a universal tree T [CDF⁺19]. Our goal is to compute $\mu^{\mathcal{G}_\tau^\dagger}$ in polynomial time. Nevertheless, running this method in parallel with a more efficient algorithm for computing $\mu^{\mathcal{G}_\tau^\dagger}$ is still useful in ensuring that we are not slower than the progress measure algorithm overall.

3.1. The Least Fixed Point of 1-Player Games. Let (G_τ, π) be a 1-player game for Even, and let $\mu \in \mathcal{L}$ be a node labeling such that there are no loose arcs in G_τ . In the rest of the paper, we focus on developing efficient methods for computing $\mu^{\mathcal{G}_\tau^\dagger}$. We know that

applying the operators in $\mathcal{G}_\tau^\uparrow$ to μ is not polynomial in general. So, we will approach $\mu^{\mathcal{G}_\tau^\uparrow}$ from above instead.

Given a node labeling $\nu : V \rightarrow \bar{L}(T)$ and an arc $vw \in E$, let $\text{drop}(\nu, vw)$ be the largest element $\xi \in \bar{L}(T)$ such that $\xi \leq \nu(v)$ and vw is not loose after setting $\nu(v)$ to ξ . Observe that if vw is loose, then $\text{drop}(\nu, vw)$ is given by $\text{TIGHTEN}(\nu, vw)$. Otherwise, it is equal to $\nu(v)$. Hence, it can be computed in $\gamma(T)$ time.

Example 3.2. Let $vw \in E$ be an arc where $\pi(v) = 2$. If ν is a node labeling from the perfect (3,2)-universal tree such that $\nu(v) = (1, 2)$ and $\nu(w) = (1, 1)$, then $\text{drop}(\nu, vw) = (1, 0)$. If ν' is a node labeling from the succinct (3,2)-universal tree such that $\nu'(v) = (0, \varepsilon)$ and $\nu'(w) = (\varepsilon, \varepsilon)$, then $\text{drop}(\nu', vw) = (0, \varepsilon)$.

We are ready to define the deflationary counterpart of Lift_{vw} . For every arc $vw \in E_\tau$, define the operator $\text{Drop}_{vw} : \mathcal{L} \times V \rightarrow \bar{L}(T)$ as

$$\text{Drop}_{vw}(\nu, u) := \begin{cases} \text{drop}(\nu, vw), & \text{if } u = v \\ \nu(v), & \text{otherwise.} \end{cases}$$

For a subgraph $H \subseteq G_\tau$, we denote $\mathcal{H}^\downarrow := \{\text{Drop}_e : e \in E(H)\}$ as the operators in H . Since they are deflationary and monotone, for any $\nu \in \mathcal{L}$, the *greatest* simultaneous fixed point $\nu^{\mathcal{H}^\downarrow}$ exists by Proposition 2.5 (i). Note that a node labeling is a simultaneous fixed point of \mathcal{H}^\downarrow if and only if there are no loose arcs in H with respect to it.

Our techniques are inspired by the methods of *label-correcting* and *label-setting* for the shortest path problem. In the shortest path problem, we have a designated target node t whose label is initialized to 0. For us, the role of t is replaced by a (potentially empty) set of even cycles in G_τ , which we do not know a priori. So, we define a set of candidate nodes called *base nodes*, whose labels need to be initialized properly.

Definition 3.3. Given a 1-player game (G_τ, π) for Even, we call $v \in V$ a *base node* if $v \in \Pi(C)$ for some even cycle C in G_τ . Denote $B(G_\tau)$ as the set of base nodes in G_τ .

Next, we introduce the `MINBOTTLENECKCYCLES` problem, which will play an important role in this paper. The input is a strongly connected digraph $G = (V, E)$ with arc costs $c \in \mathbb{R}^E$. For a cycle C , we define its *bottleneck cost* as $c(C) := \max_{e \in E(C)} c(e)$. The goal is to compute for every node $v \in V$, the minimum bottleneck cost of a cycle which contains v . It can be solved efficiently using a slight variation of the algorithm by [KKV01] for finding even cycles in parity automata.

Theorem 3.4. *Given a strongly connected digraph $G = (V, E)$ with n nodes, m arcs and ℓ distinct arc costs, `MINBOTTLENECKCYCLES` can be solved in $O(m \log \ell)$ time.*

Proof. The main idea is to recursively decompose the graph into strongly connected components (SCCs) using binary search. Let $c_1 < c_2 < \dots < c_\ell$ be the distinct arc costs. Consider the subgraph $H \subseteq G$ induced by arcs with cost at most $c_{\lceil \ell/2 \rceil}$, and let K_1, K_2, \dots, K_t be its SCCs. They can be computed in $O(n + m) = O(m)$ time using Tarjan's SCCs algorithm [Tar72]. For each $i \in [t]$, if $|E(K_i)| > 0$, then every node in K_i has a cycle going through it with cost at most $c_{\lceil \ell/2 \rceil}$. Otherwise, K_i is a singleton with no self-loops. Hence, every cycle going through it has cost greater than $c_{\lceil \ell/2 \rceil}$. This observation allows us to split G into $t + 1$ subinstances G_1, G_2, \dots, G_{t+1} . The first t subinstances are given by $G_i = (V_i, E_i) = K_i$ for all $i \in [t]$. The last subinstance $G_{t+1} = (V_{t+1}, E_{t+1})$ is obtained by contracting each K_i into a node in G , and destroying any self-loops on the resulting node. Then, the procedure is

repeated on these subinstances. In particular, for each $i \in [t]$, we consider the subgraph $H_i \subseteq G_i$ induced by arcs with cost at most $c_{\lceil \ell/4 \rceil}$. For G_{t+1} , we consider the subgraph $H_{t+1} \subseteq G_{t+1}$ induced by arcs with cost at most $c_{\lceil 3\ell/4 \rceil}$.

By construction, every edge in G contributes an edge to at most one of the $t + 1$ subinstances. Hence, $\sum_{i=1}^{t+1} |E_i| \leq |E|$. Furthermore, for every $i \in [t + 1]$, we have $|E_i| \geq |V_i|$ because G_i is strongly connected. It follows that the total running time of all $t + 1$ calls to Tarjan's SCCs algorithm is $O(m)$ again. Since the recursion depth is $O(\log \ell)$, the overall running time is $O(m \log \ell)$. \square

Computing the set of base nodes in G_τ can be reduced to MINBOTTLENECKCYCLES. For every node $v \in V$, assign cost $c(e) := \pi(v)$ to all its outgoing arcs $e \in \delta_{G_\tau}^+(v)$. Then, apply Theorem 3.4 to every SCC of G_τ which contains at least one arc. It is easy to see that $v \in B(G_\tau)$ if and only if $\pi(v)$ is even and the minimum bottleneck cost of a cycle passing through v is $\pi(v)$. As this reduction takes $O(n + m) = O(m)$ time, computing $B(G_\tau)$ takes $O(m \log d)$ time.

We start by presenting a label-setting method for computing $\mu^{\mathcal{G}_\tau^\dagger}$, and demonstrate its applicability to perfect universal trees. Then, we develop a label-correcting method for computing $\mu^{\mathcal{G}_\tau^\dagger}$, and apply it to the quasi-polynomial universal trees constructed in the literature [JL17, DJT20].

4. LABEL-SETTING METHOD FOR COMPUTING THE LEAST FIXED POINT

The epitome of a label-setting algorithm is none other than Dijkstra's algorithm [Dij59]. In this section, we develop its analogue for ordered trees. It can be used to compute $\mu^{\mathcal{G}_\tau^\dagger}$ whenever the labels $\{\mu^{\mathcal{G}_\tau^\dagger}(v) : v \in B(G_\tau)\}$ are known beforehand. We show that these labels are easily obtained for perfect universal trees, so it can be used to compute $\mu^{\mathcal{G}_\tau^\dagger}$ for these trees.

The algorithm takes as input a 1-player game (G_τ, π) for Even and a node labeling $\nu : V \rightarrow \bar{L}(T)$ from some ordered tree T . During its execution, the node labeling ν is updated. The algorithm also maintains a growing node set $S \subseteq V$ such that $\nu(v)$ remains fixed for all $v \in S$. For the sake of brevity, let us denote $H := G_\tau \setminus \cup_{v \in B(G_\tau)} \delta^+(v)$. Observe that there are no even cycles in H .

In every iteration, a new node is added to S , whose label is fixed. To determine this node, we introduce a label function Φ , which remains fixed throughout the algorithm. It encodes a family of topological orders in H induced by the even priorities. The next node to be added to S is then selected using a *potential* function Φ^ν , defined based on the labels Φ and ν . We remark that the usual criterion of selecting a node $v \in V \setminus S$ based on the smallest label $\nu(v)$ does not work. This is because the representation of a parity game as a mean payoff game can have negative arc weights, and the correctness of Dijkstra's algorithm is not guaranteed in the presence of negative arc weights.

To describe Φ , we define a family of functions Φ_p , parametrized by the even priorities in H . For an even $p \in [d]$, the function Φ_p encodes the topological order of nodes in the subgraph H_p . Recall that H_p is the subgraph of H induced by the nodes with priority at most p . Formally, $\Phi_p : V \rightarrow \mathbb{Z}_+$ is any function which satisfies the following three properties:

- $\Phi_p(v) = 0$ if and only if $\pi(v) > p$;
- $\Phi_p(v) \geq \Phi_p(w)$ if v can reach w in H_p ;

- $\Phi_p(v) = \Phi_p(w) > 0$ if and only if v and w are strongly connected in H_p .

The label function $\Phi : V \rightarrow \mathbb{Z}_+^{d/2}$ is then defined as

$$\Phi(v) := (\Phi_d(v), \Phi_{d-2}(v), \dots, \Phi_2(v)).$$

A linear order on $\Phi(v)$ is obtained by extending the linear order of its components lexicographically.

Remark 4.1. Given a pair of nodes v and w , comparing $\Phi(v)$ and $\Phi(w)$ amounts to finding the largest $p \in [d]$ such that $\Phi_p(v) \neq \Phi_p(w)$. Observe that if $\Phi_q(v) = \Phi_q(w) > 0$ for some $q \in [d]$, then $\Phi_r(v) = \Phi_r(w)$ for all $r \geq q$. On the other hand, if $\Phi_q(v) = \Phi_q(w) = 0$, then $\Phi_r(v) = \Phi_r(w) = 0$ for all $r \leq q$. Hence, such a p can be computed in $O(\log d)$ time via binary search.

Given a node labeling $\nu : V \rightarrow \bar{L}(T)$, the potential function $\Phi^\nu : V \rightarrow (\mathbb{Z}_+^{d/2} \times L(T)) \cup \{\infty\}$ is obtained by interlacing the components of Φ and ν in the following way

$$\Phi^\nu(v) := \begin{cases} (\Phi_d(v), \nu(v)_{d-1}, \dots, \Phi_2(v), \nu(v)_1), & \text{if } \nu(v) \neq \top. \\ \infty & \text{otherwise.} \end{cases}$$

A linear order on $\Phi^\nu(v)$ is acquired by extending the linear order of its components lexicographically. For any $p \in [d]$, the p -truncation of $\Phi(v)$ and $\Phi^\nu(v)$, denoted $\Phi(v)|_p$ and $\Phi^\nu(v)|_p$ respectively, are obtained by deleting the components with index less than p , with the convention $\infty|_p := \infty$.

We are ready to state Dijkstra's algorithm for ordered trees (Algorithm 2). First, it initializes the node set S as $B(G_\tau)$, and sets $\nu(v) := \top$ for all $v \in V \setminus S$. Then, for each even $p \in [d]$, it computes the topological order Φ_p by running Tarjan's SCCs algorithm on H_p . Next, ν is updated by dropping the tail labels of the incoming arcs $\delta^-(S)$. At the start of every iteration, the algorithm selects a node u with minimum potential $\Phi^\nu(u)$ among all the nodes in $V \setminus S$ (ties are broken arbitrarily). Then, it adds u to S and updates ν by dropping the tail labels of the incoming arcs $\delta^-(u) \cap \delta^-(S)$. The algorithm terminates when $S = V$.

Algorithm 2 Dijkstra: (G_τ, π) 1-player game for Even, $\nu : V \rightarrow \bar{L}(T)$ node labeling from an ordered tree T

```

1: procedure DIJKSTRA( $(G_\tau, \pi), \nu$ )
2:    $S \leftarrow B(G_\tau)$ 
3:    $\nu(v) \leftarrow \top$  for all  $v \in V \setminus S$ 
4:   Compute topological order  $\Phi_p$  for all even  $p \in [d]$   $\triangleright$  Using Tarjan's SCCs algorithm
5:   for all  $vw \in \delta^-(S)$  do
6:      $\nu(v) \leftarrow \text{drop}(\nu, vw)$ 
7:   while  $S \subsetneq V$  do
8:      $u \in \arg \min_{v \in V \setminus S} \Phi^\nu(v)$   $\triangleright$  Break ties arbitrarily
9:      $S \leftarrow S \cup \{u\}$ 
10:    for all  $vu \in \delta^-(u)$  where  $v \notin S$  do
11:       $\nu(v) \leftarrow \text{drop}(\nu, vu)$ 
12:  return  $\nu$ 

```

An efficient implementation of Dijkstra's algorithm using Fibonacci heaps was given by Fredman and Tarjan [FT87]. Its running time is $O(m + n \log n)$ when the keys in the

heap are real numbers, assuming that each elementary operation on the reals takes constant time. In our setting, the keys are the node potentials Φ^ν . Since computing $\text{drop}(\nu, e)$ takes $\gamma(T)$ time while comparing the potential of two nodes takes $\gamma(T) + \log d$ time, their result translates to $O(\gamma(T)m + (\gamma(T) + \log d)n \log n)$ time here. Since identifying the base nodes $B(G_\tau)$ takes $O(m \log d)$ time while computing the topological orders Φ_p takes $O(dm)$ time, the total running time of Algorithm 2 is $O((\gamma(T) + d)m + (\gamma(T) + \log d)n \log n)$.

The following lemma is a consequence of no even cycles in H .

Lemma 4.2. *Let $\nu \in \mathcal{L}$ be a node labeling. For every non-violated arc $vw \in E(H)$, we have $\Phi^\nu(v) \geq \Phi^\nu(w)$.*

Proof. We may assume that $\nu(v) \neq \top$, as otherwise $\Phi^\nu(v) = \infty$. Since v can reach w in H , we have $\Phi(v)|_{\pi(v)} \geq \Phi(w)|_{\pi(v)}$. We also have $\nu(v)|_{\pi(v)} \geq \nu(w)|_{\pi(v)}$ because vw is non-violated. Combining these two inequalities yields $\Phi^\nu(v)|_{\pi(v)} \geq \Phi^\nu(w)|_{\pi(v)}$.

For the purpose of contradiction, suppose that $\Phi^\nu(v) < \Phi^\nu(w)$. Then, $\Phi^\nu(v)|_{\pi(v)} = \Phi^\nu(w)|_{\pi(v)}$. Since this implies $\nu(v)|_{\pi(v)} = \nu(w)|_{\pi(v)}$, we conclude that $\pi(v)$ is even because vw is non-violated and $\nu(v), \nu(w) \neq \top$. It follows that $0 < \Phi_{\pi(v)}(v) = \Phi_{\pi(v)}(w)$; recall that Φ is indexed by the even priorities. Hence, v and w are strongly connected in $H_{\pi(v)}$, so there exists a cycle C in $H_{\pi(v)}$ such that $vw \in E(C)$. As v has the largest priority in $H_{\pi(v)}$, we have $\pi(C) = \pi(v)$. Thus, C is even, which is a contradiction. \square

The next theorem shows that Algorithm 2 returns the pointwise minimal node labeling which is feasible in H . The key observation is that the sequence of node potentials admitted to S during the algorithm is monotonically nondecreasing.

Theorem 4.3. *Given an initial node labeling $\nu \in \mathcal{L}$, Algorithm 2 returns the pointwise minimal node labeling $\nu^* \in \mathcal{L}$ which is feasible in H and satisfies $\nu^*(v) = \nu(v)$ for all $v \in B(G_\tau)$.*

Proof. Let $\nu_0 := \nu$ be the input node labeling, and suppose that the algorithm ran for $k - 1$ iterations. For every $i \geq 1$, let ν_i be the node labeling at the start of iteration i . Then, the algorithm returns ν_k . Clearly, ν_0 is feasible in H because $\nu_0(v) = \top$ for all $v \notin B(G_\tau)$. Since ν_i is obtained from ν_{i-1} by a sequence of $\text{drop}(\cdot, \cdot)$ operations, and this operation preserves feasibility, it follows that ν_k is feasible in H . Furthermore, $\nu_k(v) = \nu_{k-1}(v) = \dots = \nu_0(v)$ for all $v \in B(G_\tau)$ due to our initialization $S \leftarrow B(G_\tau)$.

For every $i \geq 1$, let u_i be the node added to S in iteration i . We start by showing that $\Phi^{\nu_k}(u_i) \geq \Phi^{\nu_k}(u_{i-1})$ for all $i > 1$. This is equivalent to $\Phi^{\nu_i}(u_i) \geq \Phi^{\nu_i}(u_{i-1})$ because $\nu_k(u_i) = \nu_j(u_i)$ for all $j \geq i$. If $\nu_i(u_i) = \nu_{i-1}(u_i)$, then we are done by our choice of u_{i-1} . So, we may assume that $\nu_i(u_i) < \nu_{i-1}(u_i)$, which implies that $\text{drop}(\cdot, \cdot)$ was applied to $u_i u_{i-1} \in E(H)$ in iteration $i - 1$. Hence, the arc $u_i u_{i-1}$ is tight with respect to ν_i . By Lemma 4.2, we obtain $\Phi^{\nu_i}(u_i) \geq \Phi^{\nu_i}(u_{i-1})$.

Next, we show that there are no loose arcs in H with respect to ν_k . Since $B(G_\tau)$ is an independent set in H , it suffices to prove that $u_i u_j$ is not loose with respect to ν_k for all $1 \leq i < j$ where $u_i u_j \in E(H)$. For the purpose of contradiction, let $u_i u_j \in E(H)$ be a loose arc with respect to ν_k . Note that $\nu_k(u_i) \neq \top$, as otherwise it would imply $\nu_k(u_j) = \top$ because $\Phi^{\nu_k}(u_i) \leq \Phi^{\nu_k}(u_j)$. Since the label $\nu_k(u_i)$ was given by TIGHTEN, it is the smallest leaf in the subtree of T rooted at $\nu_k(u_i)|_{\pi(u_i)}$. Therefore, $\nu_k(u_i)|_{\pi(u_i)} > \nu_k(u_j)|_{\pi(u_i)}$. Let p be the smallest even integer such that $p \geq \pi(u_i)$. Then, $\Phi(u_i)|_p \geq \Phi(u_j)|_p$ because either

$\pi(u_j) > p$ or $u_i u_j \in E(H_p)$. However, these two inequalities yield $\Phi^{\nu_k}(u_i) > \Phi^{\nu_k}(u_j)$, which is a contradiction.

It is left to show the pointwise minimality of ν_k . Let $\nu' : V \rightarrow \bar{L}(T)$ be a node labeling feasible in H such that $\nu'(v) = \nu(v)$ for all $v \in B(G_\tau)$. We prove that $\nu' \geq \nu_k$. Let $u \in V$; we may assume that $\nu'(u) \neq \top$. From the definition of feasibility, there exists a strategy σ for Even such that $G_{\sigma\tau} \setminus \cup_{v \in B(G_\tau)} \delta^+(v)$ does not contain violated arcs with respect to ν' . Observe that u can reach $B(G_\tau)$ in this subgraph. Indeed, if u cannot reach $B(G_\tau)$, then it reaches a cycle C . As there are no even cycles in H , C is odd and so $\nu'(v) = \top$ for all $v \in V(C)$. Now, let P be a u - w path in this subgraph for some $w \in B(G_\tau)$. Since P has no loose arcs with respect to ν_k and $\nu_k(w) = \nu(w) = \nu'(w)$, an inductive argument from w along P shows that $\nu_k(v) \leq \nu'(v)$ for all $v \in V(P)$. \square

Consequently, if we can determine $\mu^{\mathcal{G}_\tau^\uparrow}(v)$ for all $v \in B(G_\tau)$ beforehand, then we can use Algorithm 2 to compute $\mu^{\mathcal{G}_\tau^\uparrow}$.

Corollary 4.4. *Let $\mu \in \mathcal{L}$ be a node labeling with no loose arcs in G_τ . Given an initial node labeling $\nu \in \mathcal{L}$ where $\nu(v) = \mu^{\mathcal{G}_\tau^\uparrow}(v)$ for all $v \in B(G_\tau)$, Algorithm 2 returns $\mu^{\mathcal{G}_\tau^\uparrow}$.*

Proof. Let ν^* be the node labeling returned by Algorithm 2. By Theorem 4.3, we have $\nu^* \leq \mu^{\mathcal{G}_\tau^\uparrow}$ because H is a subgraph of G_τ . For every $v \in B(G_\tau)$, observe that $\nu^*(v) = \mu^{\mathcal{G}_\tau^\uparrow}(v)$ and v has a non-violated outgoing arc with respect to $\mu^{\mathcal{G}_\tau^\uparrow}$. Hence, ν^* is feasible in G_τ . To show that $\nu^* \geq \mu^{\mathcal{G}_\tau^\uparrow}$, it suffices to prove that $\nu^* \geq \mu$ due to the pointwise minimality of $\mu^{\mathcal{G}_\tau^\uparrow}$. Let $u \in V$; we may assume that $\nu^*(u) \neq \top$. Let σ be a strategy for Even such that $G_{\sigma\tau}$ does not contain violated arcs with respect to ν^* . As $\nu^*(u) \neq \top$, u can reach an even cycle in $G_{\sigma\tau}$. So, there exists a u - w path P for some $w \in B(G_\tau)$ in $G_{\sigma\tau}$. Since P has no loose arcs with respect to μ and $\mu(w) \leq \mu^{\mathcal{G}_\tau^\uparrow}(w) = \nu^*(w)$, an inductive argument from w along P shows that $\mu(v) \leq \nu^*(v)$ for all $v \in V(P)$. \square

4.1. Application to Perfect Universal Trees. In this subsection, we show how to compute $\mu^{\mathcal{G}_\tau^\uparrow}$ using Algorithm 2 for perfect universal trees. Given a node labeling $\mu \in \mathcal{L}$ such that G_τ has no loose arcs, Algorithm 3 first applies the operator $\text{Lift}_v \in \mathcal{G}_\tau^\uparrow$ once for each base node $v \in B(G_\tau)$. Then, it calls Algorithm 2 with this updated node labeling.

Algorithm 3 Label-Setting for Perfect Universal Trees: (G_τ, π) 1-player game for Even, $\mu : V \rightarrow \bar{L}(T)$ node labeling from a perfect $(n, d/2)$ -universal tree T such that G_τ does not contain loose arcs.

```

1: procedure LABELSETTING( $(G_\tau, \pi), \mu$ )
2:    $\nu \leftarrow \mu$ 
3:   for all  $v \in B(G_\tau)$  do
4:      $\nu(v) \leftarrow \min_{vw \in E_\tau} \text{lift}(\nu, vw)$ 
5:    $\nu \leftarrow \text{DIJKSTRA}((G_\tau, \pi), \nu)$ 
6:   return  $\nu$ 

```

Theorem 4.5. *Algorithm 3 returns $\mu^{\mathcal{G}_\tau^\uparrow}$ in $O(d(m + n \log n))$ time.*

Proof. Since there are no loose arcs in G_τ , we have $\mu(v)|_{\pi(v)} \leq \mu(w)|_{\pi(v)}$ for all $vw \in E_\tau$ where $v \in B(G_\tau)$. Furthermore, if $\mu(v)|_{\pi(v)} = \mu(w)|_{\pi(v)}$, then $\mu(v)$ is either \top or the smallest leaf in the subtree of T rooted at $\mu(v)|_{\pi(v)}$. Let $\nu \in \mathcal{L}$ be the node labeling right before calling DIJKSTRA. It is obtained by applying the operators $\{\text{Lift}_v : v \in B(G_\tau)\} \subseteq \mathcal{G}_\tau^\dagger$ to μ . Hence, for every base node $v \in B(G_\tau)$, $\nu(v)$ is either \top or the smallest leaf in the subtree of T rooted at $\mu(v)|_{\pi(v)}$. Moreover, we maintain $\nu \leq \mu^{\mathcal{G}_\tau^\dagger}$ and the absence of loose arcs in G_τ with respect to ν .

First, we prove correctness. By Corollary 4.4, it suffices to prove that $\nu(v) = \mu^{\mathcal{G}_\tau^\dagger}(v)$ for all $v \in B(G_\tau)$. Fix a base node $w \in B(G_\tau)$. We may assume that $\nu(w) \neq \top$. Let C be a cycle dominated by w in G_τ , and consider the path $P := C \setminus wu$ where $wu \in E(C)$. Let $\bar{\nu} : V \rightarrow \bar{L}(T)$ be a node labeling such that $\bar{\nu}(v) = \top$ for all $v \notin V(P)$, $\bar{\nu}(w) = \nu(w)$ and P is tight with respect to $\bar{\nu}$. Then, $\bar{\nu}$ is feasible in $G_\tau \setminus \delta^+(w)$. Moreover, $\bar{\nu} \geq \nu$ because there are no loose arcs in P with respect to ν . Now, recall that $\nu(w)$ is the smallest leaf in the subtree of T rooted at $\nu(w)|_{\pi(w)}$. Since $\pi(v) \leq \pi(w)$ for all $v \in V(P)$, we have $\bar{\nu}(u)|_{\pi(w)} = \bar{\nu}(w)|_{\pi(w)}$ because $|V(P)| \leq n$. As $\pi(w)$ is even, the arc wu is tight with respect to $\bar{\nu}$. It follows that $\bar{\nu}$ is feasible in G_τ . Thus, $\nu \leq \mu^{\mathcal{G}_\tau^\dagger} \leq \nu^{\mathcal{G}_\tau^\dagger} \leq \bar{\nu}$, where the second inequality is due to $\mu \leq \nu$ and the third inequality is due to the pointwise minimality of $\nu^{\mathcal{G}_\tau^\dagger}$. In particular, $\nu(w) = \mu^{\mathcal{G}_\tau^\dagger}(w)$.

Next, we prove the running time. Applying the operator $\text{Lift}_v \in \mathcal{G}_\tau^\dagger$ once for each base node $v \in B(G_\tau)$ involves calling the TIGHTEN subroutine at most m times. On the perfect $(n, d/2)$ -perfect universal tree T , TIGHTEN takes $\gamma(T) = O(d)$ time. Finally, DIJKSTRA with T takes $O(d(m + n \log n))$ time. \square

Remark 4.6. The purpose of Line 4 in Algorithm 3 is to ensure that before calling DIJKSTRA, $\nu(v)$ is either \top or the smallest leaf in the subtree of T rooted at $\nu(v)|_{\pi(v)}$ for every base node $v \in B(G_\tau)$. If Algorithm 3 is used to compute $\mu^{\mathcal{G}_\tau^\dagger}$ in every iteration of Algorithm 1, then μ already satisfies this property due to the pointwise minimality of $\mu^{\mathcal{G}_\tau^\dagger}$. So, Line 4 can be skipped.

5. LABEL-CORRECTING METHOD FOR COMPUTING THE LEAST FIXED POINT

In the previous section, we have seen that knowing the labels $\{\mu^{\mathcal{G}_\tau^\dagger}(v) : v \in B(G_\tau)\}$ allows one to compute the least fixed point $\mu^{\mathcal{G}_\tau^\dagger}$ using DIJKSTRA. We have also shown how to compute these labels for perfect universal trees. For more complicated universal trees such as the quasi-polynomial ones, it is unclear how to compute these labels efficiently. In this section, we develop a label-correcting method for computing $\mu^{\mathcal{G}_\tau^\dagger}$. It has the advantage of only requiring an over-approximation of $\mu^{\mathcal{G}_\tau^\dagger}(v)$ for all $v \in B(G_\tau)$.

The Bellman–Ford algorithm for the shortest path problem is a well-known implementation of the generic label-correcting method [AMO93]. We start by giving its analogue for ordered trees. Algorithm 4 takes as input a 1-player game (G_τ, π) for Even and a node labeling $\nu : V \rightarrow \bar{L}(T)$ from some ordered tree T . Like its classical version for shortest paths, the algorithm runs for $n - 1$ iterations. In each iteration, it replaces the tail label of every arc $e \in E_\tau$ by $\text{drop}(\nu, e)$. Clearly, the running time is $O(mn\gamma(T))$. Moreover, if ν' is the returned node labeling, then $\nu' \geq \nu^{\mathcal{G}_\tau^\dagger}$.

Algorithm 4 Bellman–Ford: (G_τ, π) 1-player game for Even, $\nu : V \rightarrow \bar{L}(T)$ node labeling from an ordered tree T

```

1: procedure BELLMANFORD( $(G, \pi), \nu$ )
2:   for  $i = 1$  to  $n - 1$  do
3:     for all  $vw \in E$  do  $\triangleright$  In any order
4:        $\nu(v) \leftarrow \text{drop}(\nu, vw)$ 
5:   return  $\nu$ 

```

Recall that we have a node labeling $\mu \in \mathcal{L}$ such that G_τ does not have loose arcs, and our goal is to compute $\mu^{\mathcal{G}_\tau^\dagger}$. We now specify an upper bound of $\nu(v)$ for all $v \in B(G_\tau)$ which ensures that Algorithm 4 does not return a node labeling larger than $\mu^{\mathcal{G}_\tau^\dagger}$.

Definition 5.1. Given a node labeling $\mu \in \mathcal{L}$, define $\hat{\mu} : B(G_\tau) \rightarrow \bar{L}(T)$ as

$$\hat{\mu}(v) := \min_{\tilde{\mu} \in \mathcal{L}} \{ \tilde{\mu}(v) : \tilde{\mu}(v) \geq \mu(v) \text{ and } \tilde{\mu} \text{ is feasible in a cycle dominated by } v \text{ in } G_\tau \} .$$

We call $\hat{\mu}(v)$ the *threshold label* of base node v .

The next lemma follows directly from the pointwise minimality of $\mu^{\mathcal{G}_\tau^\dagger}$.

Lemma 5.2. Let $\mu \in \mathcal{L}$ be a node labeling such that G_τ does not have loose arcs. For every base node $v \in B(G_\tau)$, the threshold label of v satisfies $\hat{\mu}(v) \geq \mu^{\mathcal{G}_\tau^\dagger}(v)$.

Proof. Fix a base node v . Let C be a cycle dominated by v in G_τ . Let $\tilde{\mu} \in \mathcal{L}$ be a node labeling which is feasible in C and satisfies $\tilde{\mu}(v) \geq \mu(v)$. It suffices to prove that $\tilde{\mu}(v) \geq \mu^{\mathcal{G}_\tau^\dagger}(v)$, as C and $\tilde{\mu}$ were chosen arbitrarily. Without loss of generality, we may assume that $\tilde{\mu}(w) = \top$ for all $w \in V \setminus V(C)$. Then, $\tilde{\mu}$ is feasible in G_τ . Let u be the unique out-neighbour of v in C , and consider the u - v subpath P of C . Since $\tilde{\mu}(v) \geq \mu(v)$ and there are no loose arcs in P with respect to μ , an inductive argument from v along P shows that $\tilde{\mu}(w) \geq \mu(w)$ for all $w \in V(P)$. It follows that $\tilde{\mu} \geq \mu$. Hence, $\tilde{\mu}$ is a fixed point of \mathcal{G}_τ^\dagger that is pointwise at least μ . From the pointwise minimality of $\mu^{\mathcal{G}_\tau^\dagger}$, we get $\tilde{\mu} \geq \mu^{\mathcal{G}_\tau^\dagger}$. \square

Even though $\mu^{\mathcal{G}_\tau^\dagger}$ is feasible in G_τ , there may exist a base node v such that $\mu^{\mathcal{G}_\tau^\dagger}$ is infeasible in every cycle dominated by v in G_τ (see Figure 5). Hence, we do not have the reverse inequality.




Figure 5: An example of a strategy subgraph G_τ with node labelings from the succinct $(3,2)$ -universal tree (μ on the left and $\mu^{\mathcal{G}_\tau^\dagger}$ on the right). Nodes in V_0 and V_1 are drawn as squares and circles respectively. The base nodes are v, w where $\pi(v) = 2$ and $\pi(w) = 4$. Note that $\hat{\mu}(v) = (\varepsilon, 0) > \mu^{\mathcal{G}_\tau^\dagger}(v)$.

The next theorem shows that if we initialize the base nodes with their corresponding threshold labels, then Algorithm 4 returns $\mu^{\mathcal{G}_\tau^\uparrow}$. Even more, it suffices to have an initial node labeling $\nu \in \mathcal{L}$ such that $\mu^{\mathcal{G}_\tau^\uparrow}(v) \leq \nu(v) \leq \widehat{\mu}(v)$ for all $v \in B(G_\tau)$. For the other nodes $v \notin B(G_\tau)$, we can simply set $\nu(v) \leftarrow \top$.

Theorem 5.3. *Let $\mu \in \mathcal{L}$ be a node labeling such that G_τ does not have loose arcs. For every base node $v \in B(G_\tau)$, let $\widehat{\mu}(v)$ be the threshold label of v . Given an input node labeling $\nu \in \mathcal{L}$ which satisfies $\nu \geq \mu^{\mathcal{G}_\tau^\uparrow}$ and $\nu(v) \leq \widehat{\mu}(v)$ for all $v \in B(G_\tau)$, Algorithm 4 returns $\mu^{\mathcal{G}_\tau^\uparrow}$.*

Proof. Given input $\nu \in \mathcal{L}$, let ν' be the node labeling returned by Algorithm 4. Then, we have

$$\mu \leq \mu^{\mathcal{G}_\tau^\uparrow} \leq \nu^{\mathcal{G}_\tau^\downarrow} \leq \nu' \leq \nu. \quad (5.1)$$

The second inequality is justified as follows. Since G_τ does not have loose arcs with respect to μ , it also does not have loose arcs with respect to $\mu^{\mathcal{G}_\tau^\uparrow}$ due to the pointwise minimality of $\mu^{\mathcal{G}_\tau^\uparrow}$. Hence, $\mu^{\mathcal{G}_\tau^\uparrow}$ is a fixed point of $\mathcal{G}_\tau^\downarrow$. As $\mu^{\mathcal{G}_\tau^\uparrow} \leq \nu$, we get $\mu^{\mathcal{G}_\tau^\uparrow} \leq \nu^{\mathcal{G}_\tau^\downarrow}$ by the pointwise maximality of $\nu^{\mathcal{G}_\tau^\downarrow}$. The third inequality, on the other hand, follows from the fact that $\nu^{\mathcal{G}_\tau^\downarrow}$ and ν' are obtained by iterating the operators in $\mathcal{G}_\tau^\downarrow$ on ν .

First, we show that $\mu^{\mathcal{G}_\tau^\uparrow} = \nu^{\mathcal{G}_\tau^\downarrow}$. Since $\mu^{\mathcal{G}_\tau^\uparrow}$ is feasible in G_τ , there exists a strategy σ for Even such that $G_{\sigma\tau}$ has no violated arcs with respect to $\mu^{\mathcal{G}_\tau^\uparrow}$. In the subgraph $G_{\sigma\tau}$, every node has out-degree 1, so every node can reach a cycle. Pick any cycle C in $G_{\sigma\tau}$ and let $v \in \Pi(C)$. Without loss of generality, we may assume that $\mu^{\mathcal{G}_\tau^\uparrow}(v) < \top$, as otherwise $\mu^{\mathcal{G}_\tau^\uparrow}(w) = \top$ for all the nodes w which can reach v in $G_{\sigma\tau}$. Then, C is even by the Cycle Lemma, which implies that $v \in B(G_\tau)$. From the definition of threshold label, we have $\widehat{\mu}(v) \leq \mu^{\mathcal{G}_\tau^\uparrow}(v)$. As $\nu(v) \leq \widehat{\mu}(v)$, it follows from (5.1) that $\mu^{\mathcal{G}_\tau^\uparrow}(v) = \nu^{\mathcal{G}_\tau^\downarrow}(v)$. Now, let u be a node which can reach v in $G_{\sigma\tau}$, and consider the corresponding u - v path P . Since there are no loose arcs in $G_{\sigma\tau}$ with respect to $\nu^{\mathcal{G}_\tau^\downarrow}$, an inductive argument from v along P shows that $\nu^{\mathcal{G}_\tau^\downarrow}(u) \leq \mu^{\mathcal{G}_\tau^\uparrow}(u)$. As C and u were chosen arbitrarily, $\nu^{\mathcal{G}_\tau^\downarrow} \leq \mu^{\mathcal{G}_\tau^\uparrow}$ as desired.

Next, we show that $\nu^{\mathcal{G}_\tau^\downarrow} = \nu'$, or equivalently, ν' is a fixed point of $\mathcal{G}_\tau^\downarrow$. Recall that ν' is the node labeling returned by Algorithm 4. Let $T := (n-1)m_\tau$ be the total number of steps carried out by the algorithm, where $m_\tau := |E_\tau|$. For each $0 \leq t \leq T$, let $\nu_t \in \mathcal{L}$ be the node labeling at the end of step t . Note that $\nu_0 = \nu$ and $\nu_T = \nu'$.

Proposition 5.4. *For each $0 \leq t \leq T$, if an arc $uv \in E_\tau$ is loose with respect to ν_t , then there exists a path P from u to some node w such that $uv \in E(P)$, $\nu_t(w) = \nu(w)$ and ν_t is feasible in P .*

Proof. We proceed by induction on t . The base case $t = 0$ is trivially true with the path $P = uv$. Now, suppose that the lemma is true for some $t \geq 0$. Let pq be the arc processed in step $t+1$. We may assume that pq is loose with respect to ν_t , as otherwise $\nu_{t+1} = \nu_t$ and we are done. By the inductive hypothesis, there exists a path P from p to some node r such that $pq \in E(P)$, $\nu_t(r) = \nu(r)$ and ν_t is feasible in P . Pick a shortest such P . Then, $\nu_t(s) < \nu(s)$ for all $s \in V(P) \setminus \{p, r\}$.

Let uv be a loose arc with respect to ν_{t+1} . There are 2 cases:

Case 1: uv was loose with respect to ν_t . By the inductive hypothesis, there exists a path Q from u to some node w such that $uv \in E(Q)$, $\nu_t(w) = \nu(w)$ and ν_t is feasible in Q . Pick a shortest such Q . Then, $\nu_t(s) < \nu(s)$ for all $s \in V(Q) \setminus \{u, w\}$. We may assume

that $p \in V(Q) \setminus \{u\}$, as otherwise we are done. Consider the walk $W = Q_{up}P$ obtained by concatenating the u - p subpath of Q and P . Every arc in W is non-violated with respect to ν_{t+1} .

Case 2: uv was not loose with respect to ν_t . In this case, $v = p$. Consider the walk $W = uvP$ obtained by concatenating uv and P . Every arc in W is non-violated with respect to ν_{t+1} .

In both cases, we obtain a u - r walk W such that $uv \in E(W)$ and $\nu_{t+1}(r) = \nu(r)$ with no violated arcs with respect to ν_{t+1} . It remains to show that W is a path. For the purpose of contradiction, let C be the first cycle encountered while traversing W from u . Note that $p \in V(C)$. Since ν_{t+1} is feasible in C and $\nu_{t+1}(p) < \nu_t(p) \leq \top$, it follows from the Cycle Lemma that $\nu_{t+1}(s) < \top$ for all $s \in V(C)$. Furthermore, C is even. Let $x \in \Pi(C)$, which is a base node. If $x = u$, then $\widehat{\mu}(x) < \nu_{t+1}(x) \leq \nu(x) \leq \widehat{\mu}(x)$, where the strict inequality is due to $uv \in E(C)$ being loose. If $x \neq u$, then $\widehat{\mu}(x) \leq \nu_{t+1}(x) < \nu(x) \leq \widehat{\mu}(x)$, where the strict inequality is due to our choice of P and Q . Both scenarios result in a contradiction. \square

Let P be a u - w path with i arcs. Let $\bar{\nu} \in \mathcal{L}$ be a node labeling for which P is tight and $\bar{\nu}(w) = \nu(w)$. It is left to prove that $\nu_{im_\tau}(u) \leq \bar{\nu}(u)$. This is because Proposition 5.4 would imply the absence of loose arcs with respect to $\nu_T = \nu'$. We proceed by induction on $i \geq 0$. The base case $i = 0$ is trivial. Suppose that the statement is true for all paths with at most i arcs, and consider $|E(P)| = i + 1$. Let uv be the first arc of P . By the inductive hypothesis, $\nu_{im_\tau}(v) \leq \bar{\nu}(v)$. Since uv is processed in iteration $i + 1$ and $\nu_{t+1} \leq \nu_t$ for all t , we obtain $\nu_{(i+1)m_\tau}(u) \leq \bar{\nu}(u)$. \square

Our strategy for computing the node labeling ν as specified in the statement of Theorem 5.3 is to find the cycles in Definition 5.1. In particular, for every base node $v \in B(G_\tau)$, we aim to find a cycle C dominated by v in G_τ such that $\widehat{\mu}(v)$ can be extended to a node labeling that is feasible in C . To accomplish this goal, we first introduce the notion of *width* in Section 5.1, which allows us to evaluate how ‘good’ a cycle is. It is defined using chains in the poset of subtrees of T , where the partial order is given by \sqsubseteq . Then, in Section 5.2, we show how to obtain the desired cycles by computing minimum bottleneck cycles on a suitably defined auxiliary digraph.

5.1. Width from a Chain of Subtrees in T . Two ordered trees T' and T'' are said to be *distinct* if $T' \not\cong T''$ (not isomorphic in the sense of Definition 2.1). Let h be the height of our universal tree T . For $0 \leq j \leq h$, denote \mathcal{T}_j as the set of distinct (whole) subtrees rooted at the vertices of depth $h - j$ in T . For example, $\mathcal{T}_h = \{T\}$, while \mathcal{T}_0 contains the trivial tree with a single vertex. Since we assumed that all the leaves in T are at the same depth, every tree in \mathcal{T}_j has height j . We denote $\mathcal{T} = \cup_{j=0}^h \mathcal{T}_j$ as the union of all these subtrees. The sets \mathcal{T} and \mathcal{T}_j form posets with respect to the partial order \sqsubseteq . The next definition is the usual chain cover of a poset, where we additionally require that the chains form an indexed tuple instead of a set.

Definition 5.5. For $0 \leq j \leq h$, let $\mathcal{C}_j = (\mathcal{C}_j^0, \mathcal{C}_j^1, \dots, \mathcal{C}_j^\ell)$ be a tuple of chains in the poset $(\mathcal{T}_j, \sqsubseteq)$. We call \mathcal{C}_j a *cover of \mathcal{T}_j* if $\cup_{k=0}^\ell \mathcal{C}_j^k = \mathcal{T}_j$. A *cover of \mathcal{T}* is a tuple $\mathcal{C} = (\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_h)$ where \mathcal{C}_j is a cover of \mathcal{T}_j for all $0 \leq j \leq h$. We refer to \mathcal{C}_j as the *j th-subcover of \mathcal{C}* . Given a cover \mathcal{C} of \mathcal{T} , we denote the trees in the chain \mathcal{C}_j^k as $T_{0,j}^k \sqsubset T_{1,j}^k \sqsubset \dots \sqsubset T_{|\mathcal{C}_j^k|-1,j}^k$.

An example of an ordered tree with its cover is given in Figure 6. We are ready to introduce the key concept of this subsection.




Figure 6: An ordered tree T of height 3, and a cover \mathcal{C}_j of \mathcal{T}_j for all $0 < j < 3$. Recall that \mathcal{T}_j is the set of distinct subtrees of T rooted at depth $3 - j$, while \mathcal{C}_j^k is the k th chain in \mathcal{C}_j .

Definition 5.6. Let \mathcal{C} be a cover of \mathcal{T} . Let H be a subgraph of G_τ and $j = \lceil \pi(H)/2 \rceil$. For a fixed chain \mathcal{C}_j^k in \mathcal{C}_j , the k th-width of H , denoted $\alpha_{\mathcal{C}}^k(H)$, is the smallest integer $i \geq 0$ such that H admits a finite feasible node labeling $\nu : V(H) \rightarrow L(T_{i,j}^k)$. If such an i does not exist, then we set $\alpha_{\mathcal{C}}^k(H) = \infty$.

Recall that $T_{i,j}^k$ is the $(i + 1)$ -th smallest tree in the chain \mathcal{C}_j^k . We are mainly interested in the case when H is a cycle, and write $\alpha^k(H)$ whenever the cover \mathcal{C} is clear from context. Observe that the definition above requires $\nu(v) \neq \top$ for all $v \in V(H)$. Hence, an odd cycle has infinite k th-width by the Cycle Lemma. As $(\mathcal{C}_j^k, \sqsubseteq)$ is a chain, H admits a finite feasible node labeling $\nu : V(H) \rightarrow L(T_{i,j}^k)$ for all $\alpha^k(H) \leq i < |\mathcal{C}_j^k|$. The next lemma illustrates the connection between the k th-width of an even cycle and its path decomposition.

Lemma 5.7. Let \mathcal{C} be a cover of \mathcal{T} . For an even cycle C , let $\Pi(C) = \{v_1, v_2, \dots, v_\ell\}$ and $j = \pi(C)/2$. Decompose C into arc-disjoint paths P_1, P_2, \dots, P_ℓ such that each P_i ends at v_i . Then, $\alpha^k(C) = \max_{i \in [\ell]} \alpha^k(P_i)$ for all $0 \leq k < |\mathcal{C}_j|$.

Proof. Fix a chain \mathcal{C}_j^k and let $\alpha^* = \max_{i \in [\ell]} \alpha^k(P_i)$. Clearly, any node labeling which is feasible in C is also feasible in P_i for all $i \in [\ell]$. So, $\alpha^k(C) \geq \alpha^*$. Next, we prove the reverse inequality. For each $i \in [\ell]$, there exists a node labeling $\nu_i : V(P_i) \rightarrow L(T_{\alpha^*,j}^k)$ which is feasible in P_i . Since v_i has out-degree 0 in P_i , we may assume that $\nu_i(v_i) = \min L(T_{\alpha^*,j}^k)$ for all $i \in [\ell]$ without loss of generality. Let us define a new node labeling ν as follows. For $v \in V(C)$, set $\nu(v) := \nu_i(v)$ where $i \in [\ell]$ is the unique index such that $\delta_{P_i}^-(v) \neq \emptyset$. Then, ν is feasible in C because $\nu(v) \neq \top$ for all $v \in V(C)$. Hence, $\alpha^k(C) \leq \alpha^*$. \square

For a base node $v \in B(G_\tau)$, let us consider the cycles in G_τ which are dominated by v . Among these cycles, we are interested in finding one with the smallest k th-width. So, we extend the notion of k th-width to base nodes in the following way.

Definition 5.8. Let \mathcal{C} be a cover of \mathcal{T} . Let $v \in B(G_\tau)$ be a base node and $j = \pi(v)/2$. For $0 \leq k < |\mathcal{C}_j|$, the k th-width of v is defined as

$$\alpha_{\mathcal{C}}^k(v) := \min \left\{ \alpha_{\mathcal{C}}^k(C) : C \text{ is a cycle dominated by } v \text{ in } G_\tau \right\}.$$

Again, we write $\alpha^k(v)$ whenever the cover \mathcal{C} is clear from context. Observe that $T_{\alpha^k(v), \pi(v)/2}^k$ is the smallest tree in the chain $\mathcal{C}_{\pi(v)/2}^k$ for which some cycle dominated by v in G_τ admits a finite feasible node labeling.

Given a leaf $\xi \in L(T)$ and integers $i, j, k \in \mathbb{Z}_{\geq 0}$, the following subroutine locates a member of the chain \mathcal{C}_j^k in T whose leaves are at least ξ and into which $T_{i,j}^k$ is embeddable.

RAISE(ξ, i, j, k)

Given a leaf $\xi \in L(T)$ and integers $i \geq 0, j \in [h], k \geq 0$, return the smallest leaf $\xi' \in L(T)$ such that (1) $\xi' \geq \xi$; and (2) ξ' is the smallest leaf in the subtree $T_{i',j}^k$ for some $i' \geq i$. If ξ' does not exist, then return \top .

This subroutine allows us to relate the k th-width $\alpha^k(v)$ of a base node v to its threshold label $\widehat{\mu}(v)$. For $0 \leq k < |\mathcal{C}_{\pi(v)/2}|$, the element returned by RAISE($\mu(v), \alpha^k(v), \frac{\pi(v)}{2}, k$) is at least $\widehat{\mu}(v)$. Moreover, the smallest such element over all k is precisely $\widehat{\mu}(v)$.

Lemma 5.9. Let $\mu \in \mathcal{L}$ be a node labeling such that G_τ does not have loose arcs. Let $v \in B(G_\tau)$ be a base node and $j = \pi(v)/2$. If $\xi^k \in \bar{L}(T)$ is the element returned by RAISE($\mu(v), \alpha^k(v), j, k$), then

$$\widehat{\mu}(v) = \min_{0 \leq k < |\mathcal{C}_j|} \xi^k.$$

Proof. First, we prove that $\widehat{\mu}(v) \leq \xi^k$ for all $0 \leq k < |\mathcal{C}_j|$. Fix a k and assume that $\xi^k \neq \top$. Then, ξ^k is the smallest leaf in a subtree $T_{i,j}^k$ of T for some $i \geq \alpha^k(v)$. From the definition of $\alpha^k(v)$, there exists a cycle C dominated by v in G_τ which admits a finite feasible node labeling $\nu : V(C) \rightarrow L(T_{i,j}^k)$. Since $\pi(v) = 2j$ and $\nu(w) \neq \top$ for all $w \in V(C)$, we may assume that $\nu(v) = \min L(T_{i,j}^k)$. Let us define a new node labeling $\tilde{\mu} : V \rightarrow \bar{L}(T)$ as follows. If $w \in V(C)$, set $\tilde{\mu}(w)$ as the concatenation of $\xi^k|_{2j}$ and $\nu(w)$, which is a leaf in T . Otherwise, set $\tilde{\mu}(w) := \top$. Then, $\tilde{\mu}$ is feasible in C and $\tilde{\mu}(v) = \xi^k \geq \mu(v)$, where the equality is due to our assumption $\nu(v) = \min L(T_{i,j}^k)$. Hence, $\tilde{\mu}(v) \geq \widehat{\mu}(v)$ from the definition of $\widehat{\mu}(v)$.

It is left to show that $\widehat{\mu}(v) \geq \xi^k$ for some $0 \leq k < |\mathcal{C}_j|$. We may assume that $\widehat{\mu}(v) \neq \top$. Let $T_{i,j}^k$ be the subtree of T rooted at $\widehat{\mu}(v)|_{2j}$. From the definition of $\widehat{\mu}(v)$, there exists a node labeling $\tilde{\mu} \in \mathcal{L}$ and a cycle C dominated by v in G_τ such that $\tilde{\mu}$ is feasible in C and $\tilde{\mu}(v) = \widehat{\mu}(v) \geq \mu(v)$. By the Cycle Lemma, $\tilde{\mu}(v)|_{2j} = \tilde{\mu}(w)|_{2j}$ for all $w \in V(C)$. Since $\tilde{\mu}$ is feasible in C and $\tilde{\mu}(w) \neq \top$ for all $w \in V(C)$, it follows that $i \geq \alpha^k(C) \geq \alpha^k(v)$. It remains to show that $\widehat{\mu}(v)$ is the smallest leaf in the subtree of T rooted at $\widehat{\mu}(v)|_{2j}$, as this would imply $\widehat{\mu}(v)$ is at least the output of RAISE($\mu(v), i, j, k$). Suppose otherwise for a contradiction. Then, $\widehat{\mu}(v) = \mu(v)$. Let u be the unique out-neighbour of v in C , and consider the u - v subpath P of C . Since $\tilde{\mu}(v) \geq \mu(v)$ and there are no loose arcs in P with respect to μ , an inductive argument from v along P shows that $\tilde{\mu}(w) \geq \mu(w)$ for all $w \in V(P)$. It

follows that $\mu(v)|_{2j} = \tilde{\mu}(v)|_{2j} = \tilde{\mu}(u)|_{2j} \geq \mu(u)|_{2j}$, so vu is loose with respect to μ . This is a contradiction. \square

The necessary number of chains in the subcover $\mathcal{C}_{\pi(v)/2}$ can be large if T is an arbitrary ordered tree. Fortunately, the universal trees constructed in the literature admit covers with small subcovers. We prove that a succinct (n, h) -universal tree has a cover with only 1 chain per subcover in Section 5.4, whereas a succinct Strahler (n, h) -universal tree (introduced by Daviaud et al. [DJT20]) has a cover with at most $\log n$ chains per subcover in Section 5.5.

Let $\rho(T, \mathcal{C})$ denote the running time of RAISE. We provide efficient implementations of RAISE for succinct universal trees and succinct Strahler universal trees in Sections 5.4–5.5. They have the same running time as TIGHTEN, i.e., $\rho(T, \mathcal{C}) = O(\log n \log h)$.

5.2. Estimating the Width of Base Nodes. In light of the previous discussion, we can now focus on computing the k th-width of a base node $w \in B(G_\tau)$. Fix a $0 \leq k < |\mathcal{C}_{\pi(w)/2}|$. Since we ultimately need a label that lies between the fixed point $\mu^{\mathcal{G}_\tau^\uparrow}(w)$ and the threshold label $\hat{\mu}(w)$ in order to initialize Algorithm 4, it suffices to compute a ‘good’ under-estimation of $\alpha^k(w)$. In this subsection, we reduce this problem to computing a minimum bottleneck cycle in an auxiliary digraph D with nonnegative arc costs $c^k \geq 0$.

For a base node $w \in B(G_\tau)$, let K_w denote the SCC containing w in $(G_\tau)_{\pi(w)}$, the subgraph of G_τ induced by nodes with priority at most $\pi(w)$. Note that $K_v = K_w$ for all $v \in \Pi(K_w)$. Let K'_w be the subgraph of K_w after deleting $\delta^-(v)$ for all $v \in \Pi(K_w) \setminus \{w\}$. Then, we define J_w as the subgraph of K'_w induced by nodes which can reach w in K'_w . These are the nodes which can reach w in K_w without encountering an intermediate node of priority $\pi(w)$. We remark that J_w may not be an induced subgraph of K_w .




Figure 7: An example of a 1-player game (G_τ, π) for Even is given on the left, with its auxiliary digraph D on the right. Nodes in V_0 and V_1 are drawn as squares and circles respectively. Base nodes are labeled as w_1, w_2, w_3, w_4 . The light gray region is K_{w_4} , while the dark gray region is J_{w_4} .

The auxiliary digraph D is constructed as follows. Its node set is $B(G_\tau)$. For every ordered pair (v, w) of distinct base nodes, add the arc vw if $v \in \Pi(J_w)$. For every base node w , we also add the self-loop $w w$ if w has an outgoing arc in J_w . Observe that $vw \in E(D)$ if and only if (1) $v \neq w, v \in \Pi(K_w)$ and v can reach w by only seeing smaller priorities on the intermediate nodes; or (2) $v = w$ and there exists a cycle C in G_τ such that $\Pi(C) = \{w\}$. Hence, every component of D is strongly connected as its vertex set is $\Pi(K_w)$ for some

base node $w \in B(G_\tau)$ (see Figure 7 for an example). For $w \in B(G_\tau)$, we denote D_w as the component in D which contains w .

To finish the description of D , all that remains is to assign the arc costs c^k . Note that the graph structure of D is independent of k . We first state a useful lemma which relates the outputs of $\text{drop}(\cdot, \cdot)$ when applied to node labelings from comparable trees.

Lemma 5.10. *Given integers $0 \leq i_1 \leq i_2$ and $j, k \geq 0$, let f be an injective and order-preserving homomorphism from $T_{i_1, j}^k$ to $T_{i_2, j}^k$ with the extension $f(\top) := \top$. Let $\nu_1 : V \rightarrow \bar{L}(T_{i_1, j}^k)$ and $\nu_2 : V \rightarrow \bar{L}(T_{i_2, j}^k)$ be node labelings such that $f(\nu_1(u)) \geq \nu_2(u)$ for all $u \in V$. For any arc $vw \in E_\tau$, we have $f(\text{drop}(\nu_1, vw)) \geq \text{drop}(\nu_2, vw)$.*

Proof. We only need to check the label at node v . If vw is not loose with respect to ν_1 , then $\text{drop}(\nu_1, vw) = \nu_1(v)$. Hence, $f(\text{drop}(\nu_1, vw)) = f(\nu_1(v)) \geq \nu_2(v) \geq \text{drop}(\nu_2, vw)$. So, we may assume that vw is loose with respect to ν_1 . Note that this implies $\nu_1(w) < \top$, which in turn implies $\nu_2(w) \leq f(\nu_1(w)) < \top$. Let $\xi \in L(T_{i_2, j}^k)$ be the unique element such that vw is tight with respect to ν_2 after setting $\nu_2(v) = \xi$. Since $\nu_2(w)|_{\pi(v)} \leq f(\nu_1(w))|_{\pi(v)} = f(\nu_1(w)|_{\pi(v)})$, we obtain $\text{drop}(\nu_2, vw) \leq \xi \leq f(\text{drop}(\nu_1, vw))$. \square

Fix a base node $w \in B(G_\tau)$ and let $j = \pi(w)/2$. Let $\mathcal{J}_w^\downarrow = \{\text{Drop}_e : e \in E(J_w)\}$ denote the set of Drop operators in J_w . For $0 \leq i < |\mathcal{C}_j^k|$, define $\lambda_{i, w}^k : V(J_w) \rightarrow \bar{L}(T_{i, j}^k)$ as the greatest simultaneous fixed point of \mathcal{J}_w^\downarrow which satisfies $\lambda_{i, w}^k(w) = \min L(T_{i, j}^k)$. By Proposition 2.5, $\lambda_{i, w}^k$ exists and can be obtained by applying the operators in \mathcal{J}_w^\downarrow to the node labeling that assigns $\min L(T_{i, j}^k)$ to w , and \top to other nodes, until convergence.

For every arc $vw \in E(D)$, we specify a range $\underline{c}^k(vw) \leq c^k(vw) \leq \bar{c}^k(vw)$ of permissible arc costs. The lower bound $\underline{c}^k(vw)$ is the smallest index of a tree in the chain \mathcal{C}_j^k for which $\lambda_{i, w}^k$ is finite at an out-neighbour of v , i.e.,

$$\underline{c}^k(vw) := \min \left\{ i : \lambda_{i, w}^k(u) \neq \top \text{ for some } u \in N_{J_w}^+(v) \right\}. \quad (5.2)$$

We use the convention that the minimum over the empty set is ∞ . By Lemma 5.10, observe that if $\lambda_{i, w}^k(u) \neq \top$, then $\lambda_{i', w}^k(u) \neq \top$ for all $i' \geq i$.

On the other hand, the upper bound $\bar{c}^k(vw)$ is the smallest k th-width of a path from an out-neighbour of v to w in J_w , i.e.

$$\bar{c}^k(vw) := \min \left\{ \alpha^k(P) : P \text{ is a } u\text{-}w \text{ path in } J_w \text{ where } u \in N_{J_w}^+(v) \right\}. \quad (5.3)$$

Recall that the k th-width of a path P is the smallest index of a tree in the chain \mathcal{C}_j^k for which P has a finite feasible node labeling.

Lemma 5.11. *For every arc $vw \in E(D)$, we have $\underline{c}^k(vw) \leq \bar{c}^k(vw)$.*

Proof. We may assume that $\bar{c}^k(vw) < \infty$. Let P be a u - w path in J_w where $u \in N_{J_w}^+(v)$ and $\alpha^k(P) = \bar{c}^k(vw)$. Let $j = \pi(w)/2$. From the definition of $\alpha^k(P)$, there exists a finite node labeling $\nu : V(P) \rightarrow L(T_{\alpha^k(P), j}^k)$ which is feasible in P . Note that $\nu(s) \geq \lambda_{\alpha^k(P), w}^k(s)$ for all $s \in V(P)$ because $\nu(w) \geq \min L(T_{\alpha^k(P), j}^k) = \lambda_{\alpha^k(P), w}^k(w)$ and there are no loose arcs in P with respect to $\lambda_{\alpha^k(P), w}^k$. It follows that $\lambda_{\alpha^k(P), w}^k(u) \leq \nu(u) < \top$, which gives $\underline{c}^k(vw) \leq \alpha^k(P)$ as desired. \square

For a cycle C in D , we define its (bottleneck) c^k -cost as $c^k(C) := \max_{e \in E(C)} c^k(e)$. Note that self-loops in D are considered cycles. The next theorem allows us to obtain the desired initial node labeling ν for Algorithm 4 by computing minimum bottleneck cycles in D .

Theorem 5.12. *Let \mathcal{C} be a cover of \mathcal{T} . Let $\mu \in \mathcal{L}$ be a node labeling such that G_τ does not have loose arcs. Fix a base node $w \in B(G_\tau)$ and let $j = \pi(w)/2$. For $0 \leq k < |\mathcal{C}_j|$, let c^k be arc costs in D_w where $\underline{c}^k \leq c^k \leq \bar{c}^k$, let i^k be the minimum c^k -cost of a cycle containing w in D_w , and let ξ^k be the label returned by $\text{RAISE}(\mu(w), i^k, j, k)$.² Then,*

$$\mu^{\mathcal{G}_\tau^\dagger}(w) \leq \min_{0 \leq k < |\mathcal{C}_j|} \xi^k \leq \widehat{\mu}(w).$$

Proof. We first prove the lower bound. Fix a $0 \leq k < |\mathcal{C}_j|$ and assume that $\xi^k \neq \top$. Let C be a minimum c^k -cost cycle in D_w containing w . Denote $C = (w_1, w_2, \dots, w_\ell)$ where $w = w_1 = w_\ell$; note that these are base nodes in G_τ . For every $s \in [\ell]$, let $\nu_s : V \rightarrow \bar{L}(T)$ be the node labeling defined by $\nu_s(w_s) := \xi^k$ and $\nu_s(v) := \top$ for all $v \neq w_s$. Consider the greatest simultaneous fixed point $\nu_s^{\mathcal{J}_{w_s}^\downarrow}$, which can be obtained by applying the Drop operators in $\mathcal{J}_{w_s}^\downarrow$ to ν_s by Proposition 2.5. Since $\pi(J_{w_s}) = \pi(w_s) = j$, we have $\nu_s^{\mathcal{J}_{w_s}^\downarrow}(w_s) = \nu_s(w_s) = \xi^k$ because ξ^k is the smallest leaf in the subtree of T rooted at $\xi^k|_j$. This also implies that $\nu_s^{\mathcal{J}_{w_s}^\downarrow}(v) \geq \xi^k$ for all $v \in V$. Furthermore, $\nu_s^{\mathcal{J}_{w_s}^\downarrow}$ is feasible in $G_\tau \setminus \delta^+(w_s)$, as ν_s is feasible in $G_\tau \setminus \delta^+(w_s)$ and feasibility is maintained throughout the application of Drop operators.

Claim 5.13. For each $1 < s \leq \ell$, there exists a node $u_s \in N_{J_{w_s}}^+(w_{s-1})$ such that $\nu_s^{\mathcal{J}_{w_s}^\downarrow}(u_s)|_j = \xi^k|_j$.

Proof. Fix an s and let $i = \underline{c}^k(w_{s-1}w_s)$. From the definition of \underline{c}^k in (5.2), there exists a node $u_s \in N_{J_{w_s}}^+(w_{s-1})$ such that $\lambda_{i, w_s}^k(u_s) \neq \top$. Recall that $\lambda_{i, w_s}^k : V(J_{w_s}) \rightarrow \bar{L}(T_{i, j}^k)$ is the greatest simultaneous fixed point of $\mathcal{J}_{w_s}^\downarrow$ subject to $\lambda_{i, w_s}^k(w_s) = \min L(T_{i, j}^k)$. We will show that $\nu_s^{\mathcal{J}_{w_s}^\downarrow}(u_s)|_j = \xi^k|_j$. From our construction of ν_s , observe that $\nu_s^{\mathcal{J}_{w_s}^\downarrow} : V \rightarrow \bar{L}(T)$ is the greatest simultaneous fixed point of $\mathcal{J}_{w_s}^\downarrow$ subject to $\nu_s^{\mathcal{J}_{w_s}^\downarrow}(w_s) = \xi^k$. Since ξ^k is returned by $\text{RAISE}(\mu(w), i^k, j, k)$, it is the smallest leaf in a copy of $T_{i', j}^k$ in the main tree T for some $i' \geq i^k$. As $c^k(C) = i^k$ and $w_{s-1}w_s \in E(C)$, we have $i^k \geq c^k(w_{s-1}w_s) \geq \underline{c}^k(w_{s-1}w_s) = i$. So, $i' \geq i$ and $T_{i', j}^k \sqsubseteq T_{i, j}^k$. By Lemma 5.10, $\lambda_{i', w_s}^k(u_s) \neq \top$, which implies $\nu_s^{\mathcal{J}_{w_s}^\downarrow}(u_s)|_j = \xi^k|_j$. \square

Now, consider the node labeling ν defined by $\nu(v) := \min_{s \in [\ell]} \nu_s^{\mathcal{J}_{w_s}^\downarrow}(v)$ for all $v \in V$. Note that $\nu(w_s) = \xi^k$ for all $s \in [\ell]$. It suffices to show that ν is feasible in G_τ and $\nu \geq \mu$. This is because it would imply $\nu \geq \mu^{\mathcal{G}_\tau^\dagger}$ by the pointwise minimality of $\mu^{\mathcal{G}_\tau^\dagger}$. In particular, $\xi^k = \nu(w) \geq \mu^{\mathcal{G}_\tau^\dagger}(w)$.

We first prove that ν is feasible in G_τ . Since $\nu_s^{\mathcal{J}_{w_s}^\downarrow}$ is feasible in $G_\tau \setminus \delta^+(w_s)$ for all $s \in [\ell]$, ν is feasible in $G_\tau \setminus \cup_{s=1}^\ell \delta^+(w_s)$. This is because for each node $u \in V \setminus \{w_1, w_2, \dots, w_\ell\}$, $\nu(u) = \nu_s^{\mathcal{J}_{w_s}^\downarrow}(u)$ for some $s \in [\ell]$ and furthermore, there exists an arc $uv \in \delta_{G_\tau}^+(u)$ which is non-violated with respect to $\nu_s^{\mathcal{J}_{w_s}^\downarrow}$. Since $\nu_s^{\mathcal{J}_{w_s}^\downarrow}(v) \geq \nu(v)$, uv is also non-violated with respect to ν . It is left to show that w_s has a non-violated outgoing arc in G_τ with respect

²We set $\xi^k = \top$ if $i^k = \infty$.

to ν for all $s \in [\ell]$. Fix $1 < s \leq \ell$. By Claim 5.13, there exists a node $u_s \in N_{J_{w_s}}^+(w_{s-1})$ such that $\nu_s^{\mathcal{J}_{w_s}^\downarrow}(u_s)|_j = \xi^k|_j$. Hence, $\nu(u_s)|_j = \xi^k|_j$. As $\nu(w_{s-1}) = \xi^k$ and $\pi(w_{s-1}) = j$, the arc $w_{s-1}u_s$ is non-violated with respect to ν .

Next, we show that $\nu_s^{\mathcal{J}_{w_s}^\downarrow} \geq \mu$ for all $s \in [\ell]$, which would imply $\nu \geq \mu$ as desired. We proceed by induction on s . For the base case $s = \ell$, we know that $\nu_\ell^{\mathcal{J}_{w_\ell}^\downarrow}$ is the greatest simultaneous fixed point of $\mathcal{J}_{w_\ell}^\downarrow$ subject to $\nu_\ell^{\mathcal{J}_{w_\ell}^\downarrow}(w_\ell) = \xi^k$. Observe that μ is also a simultaneous fixed point of $\mathcal{J}_{w_\ell}^\downarrow$ because there are no loose arcs in G_τ with respect to μ . Since $\nu_\ell^{\mathcal{J}_{w_\ell}^\downarrow}(w_\ell) = \xi^k \geq \mu(w_\ell)$, where the inequality is due to $w_\ell = w$, we get $\nu_\ell^{\mathcal{J}_{w_\ell}^\downarrow} \geq \mu$. For the inductive step, suppose that $\nu_s^{\mathcal{J}_{w_s}^\downarrow} \geq \mu$ for some $1 < s \leq \ell$. By Claim 5.13, there exists a node $u_s \in N_{J_{w_s}}^+(w_{s-1})$ such that $\nu_s^{\mathcal{J}_{w_s}^\downarrow}(u_s)|_j = \xi^k|_j$. Then, $\mu(w_{s-1}) \leq \nu_s^{\mathcal{J}_{w_s}^\downarrow}(w_{s-1}) = \xi^k$. The inequality follows from the inductive hypothesis, whereas the equality is due to the tightness of $w_{s-1}u_s$ with respect to $\nu_s^{\mathcal{J}_{w_s}^\downarrow}$. As $\nu_{s-1}^{\mathcal{J}_{w_{s-1}}^\downarrow}$ is the greatest simultaneous fixed point of $\mathcal{J}_{w_{s-1}}^\downarrow$ subject to $\nu_{s-1}^{\mathcal{J}_{w_{s-1}}^\downarrow}(w_{s-1}) = \xi^k \geq \mu(w_{s-1})$, we obtain $\nu_{s-1}^{\mathcal{J}_{w_{s-1}}^\downarrow} \geq \mu$ because μ is also a simultaneous fixed point of $\mathcal{J}_{w_{s-1}}^\downarrow$.

Lastly, we prove the upper bound, i.e., $\xi^k \leq \widehat{\mu}(w)$ for some $0 \leq k < |\mathcal{C}_j|$. We may assume that $\widehat{\mu}(w) \neq \top$. By Lemma 5.9, $\widehat{\mu}(w)$ is returned by $\text{RAISE}(\mu(w), \alpha^k(w), j, k)$ for some $0 \leq k < |\mathcal{C}_j|$. Let H be a cycle in G_τ such that $w \in \Pi(H)$ and $\alpha^k(H) = \alpha^k(w)$. Let $\Pi(H) = \{w_1, w_2, \dots, w_r\}$ and decompose H into arc-disjoint paths P_1, P_2, \dots, P_r such that P_s is a w_{s-1} - w_s path for all $s \in [r]$, where $w_0 := w_r$. Since each P_s lies in the subgraph J_{w_s} , we have $w_{s-1}w_s \in E(D)$ for all $s \in [r]$, and their union induces a cycle H' containing w in D . Then,

$$i^k \leq c^k(H') = \max_{s \in [r]} c^k(w_{s-1}w_s) \leq \max_{s \in [r]} \bar{c}^k(w_{s-1}w_s) \leq \max_{s \in [r]} \alpha^k(P_s) = \alpha^k(H) = \alpha^k(w).$$

The third inequality follows from the definition of \bar{c}^k in (5.3), while the second equality is due to Lemma 5.7. It follows that $\xi^k \leq \widehat{\mu}(w)$ because RAISE is monotone with respect to its second argument. \square

5.3. The Label-Correcting Algorithm. The overall algorithm for computing $\mu^{\mathcal{G}_\tau^\uparrow}$ is given in Algorithm 5. The main idea is to initialize the labels on base nodes via the recipe given in Theorem 5.12, before running Algorithm 4. The labels on $V \setminus B(G_\tau)$ are initialized to \top . The auxiliary graph D serves as a condensed representation of the ‘best’ paths between base nodes. The arc costs are chosen such that minimum bottleneck cycles in D give a good estimate on the width of base nodes.

In the next paragraph, we elaborate on how the arc costs c^k are computed.

Algorithm 5 Label-Correcting: (G_τ, π) 1-player game for Even, \mathcal{C} cover of \mathcal{T} for some universal tree T , $\mu : V \rightarrow \bar{L}(T)$ node labeling such that G_τ does not contain loose arcs.

```

1: procedure LABELCORRECTING( $(G_\tau, \pi), \mathcal{C}, \mu$ )
2:    $\nu(v) \leftarrow \top$  for all  $v \in V$ 
3:   Construct auxiliary digraph  $D$ 
4:   for all components  $H$  in  $D$  do
5:     for  $k = 0$  to  $|\mathcal{C}_{\pi(H)/2}| - 1$  do
6:       Assign arc costs  $c^k$  to  $H$  where  $\underline{c}^k \leq c^k \leq \bar{c}^k$ 
7:       for all  $w \in V(H)$  do
8:          $i^k \leftarrow$  minimum  $c^k$ -cost of a cycle containing  $w$  in  $H$ 
9:         if  $i^k < \infty$  then
10:           $\nu(w) \leftarrow \min(\nu(w), \text{RAISE}(\mu(w), i^k, \frac{\pi(H)}{2}, k))$ 
11:    $\nu \leftarrow \text{BELLMANFORD}((G_\tau, \pi), \nu)$ 
12:   return  $\nu$ 

```

Computing arc costs. Fix a base node $w \in B(G_\tau)$ and let $j = \pi(w)/2$. We show how to compute $c^k(vw)$ for all incoming arcs $vw \in \delta_D^-(w)$ by running BELLMANFORD on J_w . For every $0 \leq i < |\mathcal{C}_j^k|$, let $\nu_i : V(J_w) \rightarrow \bar{L}(T_{i,j}^k)$ be the node labeling given by $\nu_i(w) := \min L(T_{i,j}^k)$ and $\nu_i(u) := \top$ for all $u \neq w$. Let ν'_i be the node labeling obtained by running BELLMANFORD on J_w with ν_i . Then, the cost of $vw \in \delta_D^-(w)$ is set as

$$c^k(vw) := \min \{i : \nu'_i(u) \neq \top \text{ for some } u \in N_{J_w}^+(v)\}.$$

Proposition 5.14. *For every arc $vw \in \delta_D^-(w)$, the computed arc cost satisfies*

$$\underline{c}^k(vw) \leq c^k(vw) \leq \bar{c}^k(vw).$$

Proof. For every $0 \leq i < |\mathcal{C}_j^k|$, we have $\lambda_{i,w}^k \leq \nu'_i$ because $\lambda_{i,w}^k = \nu_i^{\mathcal{J}_w^\perp}$. Therefore, $\underline{c}^k(vw) \leq c^k(vw)$. Next, let P be a u - w path in J_w for some $u \in N_{J_w}^+(v)$. For every $i \geq \alpha^k(P)$, there exists a finite node labeling $\bar{\nu}_i : V(P) \rightarrow L(T_{i,j}^k)$ which is feasible on P . Without loss of generality, we may assume that $\bar{\nu}_i(w) = \min L(T_{i,j}^k)$. Since BELLMANFORD ran for $|V(J_w)| - 1$ iterations, an inductive argument from w along P shows that $\nu'_i \leq \bar{\nu}_i$. In particular, $\nu'_i(u) \neq \top$ for all $i \geq \alpha^k(P)$. It follows that $c^k(vw) \leq \bar{c}^k(vw)$. \square

Since the procedure above involves running BELLMANFORD on J_w $|\mathcal{C}_j^k|$ times, its running time is $O(mn\gamma(T)|\mathcal{C}_j^k|)$. If $|\mathcal{C}_j^k|$ is too big, then the procedure can be combined with binary search. For every node $u \in V(J_w)$, we want to determine $\min\{i : \nu'_i(u) \neq \top\}$. Fix a node $u \in V(J_w)$. We first call BELLMANFORD on J_w with $i = \lfloor |\mathcal{C}_j^k|/2 \rfloor$, i.e., the middle tree in the chain \mathcal{C}_j^k . By Lemma 5.10, if $\nu'_i(u) = \top$, then $\nu'_{i'}(u) = \top$ for all $i' \leq i$. Otherwise, $\nu'_{i'}(u) \neq \top$ for all $i' \geq i$. In both cases, we can discard half of the trees in \mathcal{C}_j^k . Then, the process is repeated until the correct tree is found. By applying this to every node $u \in V(J_w)$, we obtain the following running time.

Proposition 5.15. *For every base node $w \in B(G_\tau)$ and $0 \leq k < |\mathcal{C}_j|$ where $j = \pi(w)/2$, the arc costs $\{c^k(vw) : vw \in \delta_D^-(w)\}$ can be computed in $O(mn\gamma(T) \cdot \min\{|\mathcal{C}_j^k|, n \log |\mathcal{C}_j^k|\})$ time.*

We are ready to prove a generic bound on the running time of Algorithm 5 for an arbitrary universal tree T with an arbitrary cover \mathcal{C} .

Theorem 5.16. *In $O(mn^2\gamma(T) \cdot \max_{j,k} |\mathcal{C}_j| \min\{|\mathcal{C}_j^k|, n \log |\mathcal{C}_j^k|\} + n\rho(T, \mathcal{C}) \cdot \max_j |\mathcal{C}_j|)$ time, Algorithm 5 returns $\mu^{\mathcal{G}_\tau^\dagger}$.*

Proof. Correctness follows immediately from Theorem 5.12 and Theorem 5.3. It is left to prove the running time. Identifying the base nodes takes $O(m \log d)$ time (Section 3.1), while constructing the auxiliary digraph D takes $O(mn)$ time. Next, for every component H in D and $0 \leq k < |\mathcal{C}_{\pi(H)/2}|$, computing the arc costs c^k takes $O(|V(H)|mn\gamma(T) \cdot \min\{|\mathcal{C}_{\pi(H)/2}^k|, n \log |\mathcal{C}_{\pi(H)/2}^k|\})$ by Proposition 5.15. Then, computing the minimum c^k -cost cycles in H takes $O(|E(H)| \log |V(H)|)$ time by Theorem 3.4, and applying RAISE takes $O(|V(H)|\rho(T, \mathcal{C}))$ time. Finally, BELLMAN–FORD runs in $O(mn\gamma(T))$ time. \square

In the next two subsections, we apply Algorithm 5 to the quasi-polynomial universal trees constructed in the literature [JL17, DJT20]. It runs in time $O(mn^2 \log n \log d)$ time for a succinct $(n, d/2)$ -universal tree, and $O(mn^2 \log^3 n \log d)$ for a succinct Strahler $(n, d/2)$ -universal tree. The vertices in these trees are encoded using tuples of binary strings. For working with these tuples, we introduce the following notation.

Definition 5.17. Given a tuple $\xi = (\xi_{2h-1}, \xi_{2h-3}, \dots, \xi_1)$ of binary strings, denote $\zeta(\xi)$ as the number of leading zeroes in ξ_{2h-1} . We also define $\zeta(\top) := -1$.

For a pair of tuples ξ, ξ' of binary strings, note that if $\xi \geq \xi'$, then $\zeta(\xi) \leq \zeta(\xi')$ by the lexicographic order on tuples.

Definition 5.18. Given a tuple $\xi = (\xi_{2h-1}, \xi_{2h-3}, \dots, \xi_1)$ of binary strings and an integer $\kappa \geq 0$, let ξ^κ be the tuple obtained by deleting κ leading zeroes from ξ_{2h-1} . If $\kappa > \zeta(\xi)$, then $\xi^\kappa := \top$. We also define $\top^\kappa := \top$.

Example 5.19. If $\xi = (001, \varepsilon, 1)$, then $\zeta(\xi) = 2$, $\xi^1 = (01, \varepsilon, 1)$, $\xi^2 = (1, \varepsilon, 1)$ and $\xi^3 = \top$.

5.4. Application to Succinct Universal Trees. Let T be a succinct (n, h) -universal tree. Recall that every leaf $\xi \in L(T)$ corresponds to an h -tuple of binary strings where the number of bits fulfills $|\xi| \leq \lfloor \log n \rfloor$. First, we show that each \mathcal{T}_j has a cover of size 1. Equivalently, each $(\mathcal{T}_j, \sqsubseteq)$ is a chain.

Lemma 5.20. *There exists a cover \mathcal{C} of \mathcal{T} such that $|\mathcal{C}_j| = 1$ for all $0 \leq j \leq h$.*

Proof. Fix $0 \leq j \leq h$ and pick two vertices $r_1, r_2 \in V(T)$ at depth $h - j$. Let T_1 and T_2 be the subtrees of T rooted at r_1 and r_2 respectively. Every leaf $\xi_1 \in L(T_1)$ and $\xi_2 \in L(T_2)$ corresponds to a j -tuple of binary strings where $|\xi_1| \leq \lfloor \log n \rfloor - |r_1|$ and $|\xi_2| \leq \lfloor \log n \rfloor - |r_2|$. Without loss of generality, assume that $|r_1| \geq |r_2|$. Then, the identity map from $V(T_1)$ to $V(T_2)$ is an order-preserving and injective homomorphism. Hence, $T_1 \sqsubseteq T_2$. \square

Since each subcover \mathcal{C}_j of \mathcal{C} consists of a single chain, we write $\mathcal{C}_j = \mathcal{T}_j$ and omit the superscript k . The subtrees in \mathcal{T}_j are $T_{0,j} \sqsubseteq T_{1,j} \sqsubseteq \dots \sqsubseteq T_{\lfloor \log n \rfloor, j}$. Observe that every leaf $\xi \in L(T_{i,j})$ corresponds to a j -tuple of binary strings where $|\xi| \leq i$.

Next, we give an efficient implementation of the RAISE subroutine.

Lemma 5.21. *For a succinct $(n, d/2)$ -universal tree T with cover $\mathcal{C} = (\mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_{d/2})$, the RAISE(ξ, i, j, k) subroutine runs in $O(\log n \log d)$ time.*

Proof. We may assume that ξ is the smallest leaf in the subtree rooted at $\xi|_{2j}$. Otherwise, we can set it to the smallest leaf of the next subtree rooted at that depth using TIGHTEN (we return \top if this next subtree does not exist). Recall that the TIGHTEN subroutine for T also runs in $O(\log n \log d)$ time. It follows that $\xi_{2j-1} \in \{0 \cdots 0, \varepsilon\}$ and $\xi_q = \varepsilon$ for all odd $q < 2j - 1$. If $|\xi_{2j-1}| \geq i$, then we simply return ξ . Otherwise, let $p \in [d]$ be the smallest even integer such that $\xi|_p$ has a child bigger than $\xi|_{p-1}$ with at most $\lfloor \log n \rfloor - i$ bits. If p does not exist, then we return \top . Otherwise, $p > 2j$. Let $r = \lfloor \log n \rfloor - i - |\xi|_{p-1}|$. There are two cases.

Case 1: $r > 0$. Return

$$(\xi_{d-1}, \dots, \xi_{p-1} \underbrace{10 \cdots 0}_{r-1}, \varepsilon, \dots, \varepsilon, \underbrace{0 \cdots 0}_i, \varepsilon, \dots, \varepsilon)$$

where the string of i zeroes is at index $2j - 1$.

Case 2: $r \leq 0$. Denote $\xi_{p-1} = b_1 b_2 \cdots b_\ell$ where $b_q \in \{0, 1\}$ for all $q \in [\ell]$. Note that $\ell \geq 1$ by our choice of p . Furthermore, there exists a largest $t \in [\ell]$ such that $b_t = 0$ and $r' := r + \ell - (t - 1) \geq 0$. Then,

- if $p = 2j + 2$, return

$$(\xi_{d-1}, \dots, \xi_{p+1}, b_1 \cdots b_{t-1}, \underbrace{0 \cdots 0}_{i+r'}, \varepsilon, \dots, \varepsilon);$$

- if $p > 2j + 2$, return

$$(\xi_{d-1}, \dots, \xi_{p+1}, b_1 \cdots b_{t-1}, \underbrace{0 \cdots 0}_{r'}, \varepsilon, \dots, \varepsilon, \underbrace{0 \cdots 0}_i, \varepsilon, \dots, \varepsilon)$$

where the string of i zeroes is at index $2j - 1$. □

Due to the self-similar structure of succinct universal trees, the running time of Algorithm 5 given in Theorem 5.16 can be improved. The following lemma yields a faster method for computing arc costs for the auxiliary digraph D_τ . The key observation is that for any pair of trees in a chain \mathcal{T}_j , the smaller tree can be obtained from the larger tree by deleting vertices in decreasing lexicographic order. For example, a succinct (3,2)-universal tree can be obtained from a succinct (7,2)-universal tree by deleting vertices whose first component does not contain a leading zero (compare Figures 3 and 8).




Figure 8: The succinct (7, 2)-universal tree

Lemma 5.22. *Given integers $0 \leq i_1 \leq i_2$ and $j \geq 0$, let $\nu_1 : V \rightarrow \bar{L}(T_{i_1, j})$ and $\nu_2 : V \rightarrow \bar{L}(T_{i_2, j})$ be node labelings such that $\nu_1(u) = \nu_2(u)^{i_2 - i_1}$ for all $u \in V$. For any arc $vw \in E_\tau$ where $\pi(v) < 2j$, we have $\text{drop}(\nu_1, vw) = \text{drop}(\nu_2, vw)^{i_2 - i_1}$.*

Proof. Let $\xi_1 = \text{drop}(\nu_1, vw)$ and $\xi_2 = \text{drop}(\nu_2, vw)$. First, assume that vw is violated with respect to ν_1 . Then, $\nu_1(v) \neq \top$, which implies that $\nu_2(v) \neq \top$. In particular, $\zeta(\nu_2(v)) \geq i_2 - i_1$. We claim that vw is also violated with respect to ν_2 . This is clear if $\nu_1(w) \neq \top$. If $\nu_1(w) = \top$, then $\zeta(\nu_2(w)) < i_2 - i_1$. As $\nu_2(v)|_{2j-1} < \nu_2(w)|_{2j-1}$ and $\pi(v) < 2j$, the arc vw is indeed violated with respect to ν_2 . Hence, $\xi_1 = \nu_1(v)$ and $\xi_2 = \nu_2(v)$.

Next, assume that vw is not violated with respect to ν_1 . If $\nu_1(w) \neq \top$, then vw is also not violated with respect to ν_2 . It is easy to verify that $\xi_1 = \xi_2^{i_2 - i_1}$. On the other hand, if $\nu_1(w) = \top$, then $\nu_1(v) = \top$. So, we have $\zeta(\nu_2(v)) < i_2 - i_1$ and $\zeta(\nu_2(w)) < i_2 - i_1$. Since $\pi(v) < 2j$, we also have $\zeta(\xi_2) < i_2 - i_1$. Thus, $\xi_1 = \top = \xi_2^{i_2 - i_1}$ as required. \square

Computing arc costs. Fix a component H in the auxiliary graph D and let $j = \pi(H)/2$. Fix a node $w \in V(H)$. We first recall the generic method for computing the arc costs of $\delta_D^-(w)$. For each $0 \leq i \leq \lfloor \log n \rfloor$, let $\nu_i : V(J_w) \rightarrow \bar{L}(T_{i, j})$ be the node labeling defined by $\nu_i(w) := \min L(T_{i, j})$ and $\nu_i(u) := \top$ for all $u \neq w$. Let ν'_i be the node labeling obtained by running BELLMANFORD on J_w with input node labeling ν_i . Then, the cost of every arc $vw \in \delta_D^-(w)$ is set to

$$c(vw) := \min \{i : \nu'_i(u) \neq \top \text{ for some } u \in N_{J_w}^+(v)\}.$$

The above method involves running BELLMANFORD $\lfloor \log n \rfloor + 1$ times. We outline a faster method which only runs BELLMANFORD once. Observe that for any $0 \leq i \leq \lfloor \log n \rfloor$, we have $\nu_i(u) = \nu_{\lfloor \log n \rfloor}(u)^{\lfloor \log n \rfloor - i}$ for all $u \in V$. Since $\delta_{J_w}^-(v) = \emptyset$ for all $v \in \Pi(J_w) \setminus \{w\}$ by construction, and $\nu_i(w) = \min L(T_{i, j})$, Lemma 5.22 yields

$$\nu'_i(u) = \nu'_{\lfloor \log n \rfloor}(u)^{\lfloor \log n \rfloor - i}$$

for all $u \in V(J_w) \setminus (\Pi(J_w) \setminus \{w\})$. In other words, from $\nu'_{\lfloor \log n \rfloor}$ we can infer $\nu'_i(u)$ for all $u \in V(J_w) \setminus (\Pi(J_w) \setminus \{w\})$ and $0 \leq i \leq \lfloor \log n \rfloor$. So, we only need to run BELLMANFORD on J_w once to obtain $\nu'_{\lfloor \log n \rfloor}$. Note that $\nu'_{\lfloor \log n \rfloor}$ is finite because $T_{\lfloor \log n \rfloor, j}$ is an (n, j) -universal tree and every node in J_w can reach w . Then, the cost of each arc $vw \in \delta_D^-(w)$ is given by

$$c(vw) = \lfloor \log n \rfloor - \max_{u \in N_{J_w}^+(v)} \zeta(\nu(u)).$$

The above argument shows that we can compute c for the component H in time $O(|V(H)|mn\gamma(T))$, saving a factor of $\log n$. We are ready to prove the running time of Algorithm 5 for succinct universal trees.

Theorem 5.23. *For a succinct $(n, d/2)$ -universal tree T with cover $\mathcal{C} = (\mathcal{T}_0, \mathcal{T}_1, \dots, \mathcal{T}_{d/2})$, Algorithm 5 runs in $O(mn^2 \log n \log d)$ time.*

Proof. By Lemma 5.20, we have $|\mathcal{C}_j| = 1$ for all $0 \leq j \leq h$. Computing arc costs for the auxiliary digraph D_τ takes $O(mn^2\gamma(T))$ time as discussed above. Hence, the running time of Algorithm 5 becomes $O(n\rho(T, \mathcal{C}) + mn^2\gamma(T))$. The result then follows from $\gamma(T) = O(\log n \log d) = \rho(T, \mathcal{C})$, where the latter equality is due to Lemma 5.21. \square

5.5. Application to Succinct Strahler Universal Trees. In this subsection, we apply Algorithm 5 to succinct Strahler universal trees. Let us start by introducing the necessary definitions. The *Strahler number* of a rooted tree T is the largest height of a perfect binary tree that is a minor of T . For example, a perfect (ℓ, h) -universal tree has Strahler number 0 if $\ell = 1$, and h otherwise.

Definition 5.24. A g -Strahler (ℓ, h) -universal tree is an ordered tree T' of height h such that $T \sqsubseteq T'$ for every ordered tree T of height h , Strahler number at most g , and with at most ℓ leaves, all at depth exactly h .

In the definition above, we may assume that $g \leq \min(h, \lfloor \log \ell \rfloor)$. Daviaud et al. [DJT20] constructed a g -Strahler (ℓ, h) -universal tree with $\ell^{O(1)}(h/g)^g$ leaves. Note that this is quasipolynomial in ℓ and h because $g \leq \min(h, \lfloor \log \ell \rfloor)$. In this paper, we refer to it as a *succinct g -Strahler (ℓ, h) -universal tree*. Every leaf $\xi = (\xi_{2h-1}, \xi_{2h-3}, \dots, \xi_1)$ in this tree corresponds to an h -tuple of binary strings which satisfies the following three properties:

- (1) There are g nonempty bit strings, i.e. $|\{i : \xi_i \neq \varepsilon\}| = g$;
- (2) The total number of bits $|\xi|$ is at most $g + \lfloor \log \ell \rfloor$;
- (3) For each odd $i \in [2h]$,
 - (a) If there are $f < g$ nonempty bit strings in $\xi_{|i+1}$ and $|\xi_{|i+1}| = f + \lfloor \log \ell \rfloor$, then $\xi_i = 0$.
 - (b) If $\xi_j \neq \varepsilon$ for all odd $j \in [i]$, then ξ_j starts with 0 for all odd $j \in [i]$.

This is the construction of $\mathcal{B}_{\ell, h}^k$ in [DJT20, Definition 19]. In each string ξ_i , the first bit is called the *leading bit*, while the remaining bits are the *non-leading bits*. Properties 1 and 2 imply that ξ contains exactly g leading bits and at most $\lfloor \log \ell \rfloor$ non-leading bits (see Figure 9 for an example). Observe that if $g = h$, then the tree is identical to a succinct (ℓ, h) -universal tree. Indeed, one can arrive at the encoding of Jurdziński and Lazić [JL17] by removing the leading zero in every string.




Figure 9: The succinct 1-Strahler $(7, 2)$ -universal tree

Let T be a succinct g -Strahler (ℓ, h) -universal tree. Let $v \in V(T)$ be a vertex at depth $h - j$ for some $0 \leq j \leq h$. If v has k nonempty strings and i non-leading bits, then the subtree rooted at v is a succinct $(g - k)$ -Strahler $(2^{\lfloor \log \ell \rfloor - i}, j)$ -universal tree. Note that if $i = \lfloor \log \ell \rfloor$, then the subtree has only one leaf. This fact is actually independent of k . Indeed, varying k only yields different encodings of the same tree.

The crucial fact for obtaining a small cover is that the subtrees of T with fixed height j and fixed Strahler number k form a chain. This leads to the following statement.

Lemma 5.25. *There exists a cover \mathcal{C} of \mathcal{T} such that $|\mathcal{C}_j| \leq g$ for all $0 \leq j \leq h$.*

Proof. Fix a height $0 \leq j \leq h$. For each $0 \leq k \leq g$, let \mathcal{C}_j^k be the set of succinct k -Strahler universal trees in \mathcal{T}_j . Then, $\cup_k \mathcal{C}_j^k = \mathcal{T}_j$. It is left to show that $(\mathcal{C}_j^k, \sqsubseteq)$ is a chain for all k . Fix a k and pick two vertices $r_1, r_2 \in V(T)$ at depth $h - j$ such that they each have $g - k$ nonempty bit strings. Then, they each have $g - k$ leading bits. So, the number of non-leading bits in r_1 and r_2 are $|r_1| - (g - k)$ and $|r_2| - (g - k)$ respectively. Let T_1 and T_2 be the subtrees of T rooted at r_1 and r_2 respectively. Observe that T_1 is a succinct k -Strahler $(2^{\lfloor \log n \rfloor - |r_1| + g - k}, j)$ -universal tree. Similarly, T_2 is a succinct k -Strahler $(2^{\lfloor \log n \rfloor - |r_2| + g - k}, j)$ -universal tree. Without loss of generality, assume that $|r_1| \geq |r_2|$. Then, the identity map from $V(T_1)$ to $V(T_2)$ is an order-preserving and injective homomorphism. Hence, $T_1 \sqsubseteq T_2$. \square

Let \mathcal{C} be the cover given in the proof of Lemma 5.25, i.e. $\mathcal{C}_j = (\mathcal{C}_j^0, \mathcal{C}_j^1, \dots, \mathcal{C}_j^g)$ for all $0 \leq j \leq h$.

Observation 5.26. We make a few observations about the cover \mathcal{C} :

- (1) We have $\mathcal{C}_j^k \neq \emptyset$ if and only if $\max(0, g - (h - j)) \leq k \leq \min(j, g)$.
- (2) For any j where $\mathcal{C}_j^0 \neq \emptyset$, we have $\mathcal{C}_j^0 = (T_{0,j}^0)$ where $T_{0,j}^0$ has only one leaf.
- (3) For any j, k where $\mathcal{C}_j^k \neq \emptyset$, its smallest tree $T_{0,j}^k$ has only one leaf.

The next lemma shows that the RAISE subroutine can be implemented efficiently using this cover when its second argument is positive. This will be sufficient for running Algorithm 5, as Theorem 5.28 shows.

Lemma 5.27. *For a succinct g -Strahler $(n, d/2)$ -universal tree T with cover \mathcal{C} , if $i > 0$, then the $\text{RAISE}(\xi, i, j, k)$ subroutine runs in $O(\log n \log d)$ time.*

Proof. Note that $k > 0$ by Observation 5.26. We may assume that ξ is the smallest leaf in the subtree of T rooted at $\xi|_{2j}$. Otherwise, we can set it to the smallest leaf of the next subtree rooted at that depth using TIGHTEN (we return \top if this next subtree does not exist). It is known that the TIGHTEN subroutine for T runs in $O(\log n \log d)$ time [DJT20, Lemma 22]. This gives us $\xi_{2j-1} \in \{0 \cdots 0, \varepsilon\}$ and $\xi_q \in \{0, \varepsilon\}$ for all $q < 2j - 1$. If there are at least i non-leading bits in ξ_{2j-1} and exactly k non-empty strings among ξ_{2j-1}, \dots, ξ_1 , then we simply return ξ . Otherwise, let $p > 2j$ be the smallest even integer such that $\xi|_p$ has a child bigger than $\xi|_{p-1}$ with at most $\lfloor \log n \rfloor - i$ non-leading bits and exactly x nonempty strings for some $g - k - (p/2 - j - 1) \leq x \leq g - k$. If p does not exist, then we return \top .

Our goal is to increase ξ minimally such that the last j components form a tuple from $L(T_{i',j}^k)$ for some $i' \geq i$. Let z be the number of non-empty strings in $\xi|_{p-1}$, and let y be the number of non-leading bits in $\xi|_{p-1}$. We set $s = g - k - z$, representing the discrepancy on the number of nonempty bit strings (Property 1). Note that $-1 \leq s \leq p/2 - j$ due to our choice of p . We also set $r = \lfloor \log n \rfloor - i - y$, representing the discrepancy on the number of non-leading bits (Property 2). We split the remaining analysis into cases based on the emptiness of ξ_{p-1} and the signs of r, s .

Case 1: $\xi_{p-1} = \varepsilon$. Note that $y \leq \lfloor \log n \rfloor - i$ and $z < g - k$. So, $r \geq 0$ and $s > 0$. Return

$$(\xi_{d-1}, \dots, \xi_{p+1}, \underbrace{1 \underbrace{0 \cdots 0}_r, \underbrace{0, \dots, 0}_{s-1}}_{p-1}, \varepsilon, \dots, \varepsilon, \underbrace{0 \cdots 0}_{i+1}, \underbrace{0, \dots, 0}_{k-1}, \varepsilon, \dots, \varepsilon)$$

where the string of $i + 1$ zeroes is at index $2j - 1$. Property 3b holds because there is at least one empty string among the last $p/2$ components. As all the other properties are also fulfilled by construction, this is a valid tuple encoding a leaf.

Case 2: $\xi_{p-1} \neq \varepsilon$ and $s < 0$. Note that $s = -1$ and ξ_{p-1} has a leading zero due to our choice of p . Let t be the number of non-leading bits in ξ_{p-1} . Then, $y - t \leq \lfloor \log n \rfloor - i$, which implies that $r + t \geq 0$. Return

$$(\xi_{d-1}, \dots, \xi_{p+1}, \varepsilon, \varepsilon, \dots, \varepsilon, \underbrace{0 \cdots 0}_{i+1+r+t}, \underbrace{0, \dots, 0}_{k-1}, \varepsilon, \dots, \varepsilon)$$

where the string of $i + 1 + r + t$ zeroes is at index $2j - 1$.

Case 3: $\xi_{p-1} \neq \varepsilon$, $r > 0$ and $s \geq 0$. Return

$$(\xi_{d-1}, \dots, \xi_{p-1} \underbrace{10 \cdots 0}_{r-1}, \underbrace{0, \dots, 0}_s, \varepsilon, \dots, \varepsilon, \underbrace{0 \cdots 0}_{i+1}, \underbrace{0, \dots, 0}_{k-1}, \varepsilon, \dots, \varepsilon)$$

where the string of $i + 1$ zeroes is at index $2j - 1$. Property 3b is satisfied because the last $p/2$ components are nonempty if and only if the last $p/2$ components of ξ are nonempty.

Case 4: $\xi_{p-1} \neq \varepsilon$, $r \leq 0$ and $s \geq 0$. Note that $s \leq p/2 - j - 1$ in this case. Denote $\xi_{p-1} = b_1 b_2 \cdots b_\ell$ where $b_q \in \{0, 1\}$ for all $q \in [\ell]$. Due to our choice of p , there exists a largest $t \in [\ell]$ such that $b_t = 0$ and $r' := r + \ell - \max(t - 1, 1) \geq 0$. The following subcases remain.

- If $s < p/2 - j - 1$ and $t = 1$, return

$$(\xi_{d-1}, \dots, \xi_{p+1}, \varepsilon, \underbrace{0 \cdots 0}_{r'+1}, \underbrace{0, \dots, 0}_s, \varepsilon, \dots, \varepsilon, \underbrace{0 \cdots 0}_{i+1}, \underbrace{0, \dots, 0}_{k-1}, \varepsilon, \dots, \varepsilon)$$

where the string of $i + 1$ zeroes is at index $2j - 1$.

- If $s = p/2 - j - 1$ and $t = 1$, return

$$(\xi_{d-1}, \dots, \xi_{p+1}, \underbrace{10 \cdots 0}_{r'}, \underbrace{0, \dots, 0}_s, \underbrace{0 \cdots 0}_{i+1}, \underbrace{0, \dots, 0}_{k-1}, \varepsilon, \dots, \varepsilon)$$

where the string of $i + 1$ zeroes is at index $2j - 1$. Property 3b holds because there exists an empty string in the last $p/2$ components of ξ by our choice of p .

- If $s = 0$ and $t > 1$, return

$$(\xi_{d-1}, \dots, \xi_{p+1}, b_1 \cdots b_{t-1}, \varepsilon, \dots, \varepsilon, \underbrace{0 \cdots 0}_{i+1+r'}, \underbrace{0, \dots, 0}_{k-1}, \varepsilon, \dots, \varepsilon)$$

where the string of $i + 1 + r'$ zeroes is at index $2j - 1$.

- If $s > 0$ and $t > 1$, return

$$(\xi_{d-1}, \dots, \xi_{p+1}, b_1 \cdots b_{t-1}, \underbrace{0 \cdots 0}_{r'+1}, \underbrace{0, \dots, 0}_{s-1}, \varepsilon, \dots, \varepsilon, \underbrace{0 \cdots 0}_{i+1}, \underbrace{0, \dots, 0}_{k-1}, \varepsilon, \dots, \varepsilon)$$

where the string of $i + 1$ zeroes is at index $2j - 1$. □

We are ready to prove the running time of Algorithm 5 for succinct Strahler universal trees. Recall that for a g -Strahler $(n, d/2)$ -universal tree, we may assume that $g \leq \log n$.

Theorem 5.28. *For a succinct g -Strahler $(n, d/2)$ -universal tree T with cover \mathcal{C} , Algorithm 5 runs in $O(mn^2 g \log^2 n \log d)$ time.*

Proof. By Lemma 5.25, we have $|\mathcal{C}_j| \leq g$ for all $0 \leq j \leq h$. Furthermore, $|\mathcal{C}_j^k| \leq \log n$ for all $0 \leq j \leq h$ and $0 \leq k < |\mathcal{C}_j|$. It is known that $\text{TIGHTEN}(\mu, vw)$ takes $O(\log n \log d)$ time [DJT20, Lemma 22], whereas Lemma 5.27 shows that $\text{RAISE}(\xi, i, j, k)$ takes $O(\log n \log d)$ time when $i > 0$.

In Algorithm 5, if $\text{RAISE}(\mu(w), i^k, \pi(H)/2, k)$ is called with $i^k = 0$, then there exists a cycle C containing w in H such that $\underline{c}^k(e) \leq c^k(e) = 0$ for all $e \in E(C)$. Denote $E(C) = \{w_0 w_1, w_1 w_2, \dots, w_{r-1} w_r\}$ where $w_0 = w_r$. For each $s \in [r]$, the greatest simultaneous fixed point $\lambda_{0, w_s}^k : V(J_{w_s}) \rightarrow \bar{L}(T_{0, j}^k)$ assigns a finite label to an outneighbour of w_{s-1} . Since $T_{0, j}^k$ has only one leaf by Observation 5.26, there exists a $w_{s-1} w_s$ path in J_{w_s} which consists of only even priorities. Hence $\bar{c}^k(w_{s-1} w_s) = 0$. The same path also certifies that $\bar{c}^{k'}(w_{s-1} w_s) = \underline{c}^{k'}(w_{s-1} w_s) = 0$ for all $k' \neq k$. It follows that $c^{k'}(e) = 0$ for all $e \in E(C)$, so the minimum $c^{k'}$ -cost of a cycle containing w in H is 0 for all k' . Thus, the algorithm will always call $\text{RAISE}(\mu(w), i^{k'}, \pi(H)/2, k')$ with $i^{k'} = 0$ for all k' . As $T_{0, j}^{k'}$ has only one leaf for all k' by Observation 5.26, the minimum of $\text{RAISE}(\mu(w), 0, \pi(H)/2, k')$ over all k' is $\mu(w)$ if $\mu(w)$ is the smallest leaf of the subtree rooted at $\mu(w)|_{\pi(H)}$. Otherwise, it is the smallest leaf of the next subtree rooted at that depth, which can be obtained via TIGHTEN . The overall running time bound then follows from Theorem 5.16. \square

ACKNOWLEDGMENT

The authors are grateful for the helpful comments and support by László Végh. They would like to thank Xavier Allamigeon, Nathanaël Fijalkow and Marcin Jurdziński for inspiring discussions. They are also thankful to the anonymous reviewers for their valuable comments, which have helped improve the presentation of this paper.

REFERENCES

- [AGG12] Marianne Akian, Stéphane Gaubert, and Alexander Guterman. Tropical polyhedra are equivalent to mean payoff games. *International Journal of Algebra and Computation*, 22(01):1250001, 2012. doi:10.1142/S0218196711006674.
- [AMO93] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows. Theory, algorithms, and applications*. Englewood Cliffs, NJ: Prentice Hall, 1993.
- [BDM18] Massimo Benerecetti, Daniele Dell’Erba, and Fabio Mogavero. Solving parity games via priority promotion. *Formal Methods Syst. Des.*, 52(2):193–226, 2018. doi:10.1007/S10703-018-0315-1.
- [BDM⁺25] Massimo Benerecetti, Daniele Dell’Erba, Fabio Mogavero, Sven Schewe, and Dominik Wojtczak. Priority promotion with parysian flair. *J. Comput. Syst. Sci.*, 147:103580, 2025. doi:10.1016/J.JCSS.2024.103580.
- [BSV03] Henrik Björklund, Sven Sandberg, and Sergei G. Vorobyov. A discrete subexponential algorithm for parity games. In *20th Annual Symposium on Theoretical Aspects of Computer Science, STACS*, volume 2607 of *Lecture Notes in Computer Science*, pages 663–674, 2003. doi:10.1007/3-540-36494-3_58.
- [BV07] Henrik Björklund and Sergei G. Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. *Discret. Appl. Math.*, 155(2):210–229, 2007. doi:10.1016/J.DAM.2006.04.029.
- [CDF⁺19] Wojciech Czerwinski, Laure Daviaud, Nathanaël Fijalkow, Marcin Jurdziński, Ranko Lazic, and Pawel Parys. Universal trees grow inside separating automata: Quasi-polynomial lower bounds for parity games. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 2333–2349, 2019. doi:10.1137/1.9781611975482.142.

- [CJK⁺17] Cristian S. Calude, Sanjay Jain, Bakhadyr Khoussainov, Wei Li, and Frank Stephan. Deciding parity games in quasipolynomial time. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC*, pages 252–263, 2017. doi:10.1145/3055399.3055409.
- [Dij59] Edsger W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959. doi:10.1007/BF01386390.
- [DJT20] Laure Daviaud, Marcin Jurdzinski, and K. S. Thejaswini. The strahler number of a parity game. In *47th International Colloquium on Automata, Languages, and Programming, ICALP*, volume 168 of *LIPICs*, pages 123:1–123:19, 2020. doi:10.4230/LIPICs.ICALP.2020.123.
- [EJ91] E. Allen Emerson and Charanjit S. Jutla. Tree automata, mu-calculus and determinacy. In *32nd Annual Symposium on Foundations of Computer Science, FOCS*, pages 368–377, 1991. doi:10.1109/SFCS.1991.185392.
- [EJS93] E. Allen Emerson, Charanjit S. Jutla, and A. Prasad Sistla. On model-checking for fragments of μ -calculus. In *5th International Conference on Computer-Aided Verification, CAV*, volume 697 of *Lecture Notes in Computer Science*, pages 385–396, 1993. doi:10.1007/3-540-56922-7_32.
- [EM79] Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *Int. J. Game Theory*, 8:109–113, 1979. doi:10.1007/BF01768705.
- [FJdK⁺19] John Fearnley, Sanjay Jain, Bart de Keijzer, Sven Schewe, Frank Stephan, and Dominik Wojtczak. An ordered approach to solving parity games in quasi-polynomial time and quasi-linear space. *Int. J. Softw. Tools Technol. Transf.*, 21(3):325–349, 2019. doi:10.1007/S10009-019-00509-3.
- [Fri09] Oliver Friedmann. An exponential lower bound for the parity game strategy improvement algorithm as we know it. In *Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS*, pages 145–156, 2009. doi:10.1109/LICS.2009.27.
- [Fri11] Oliver Friedmann. *Exponential Lower Bounds for Solving Infinitary Payoff Games and Linear Programs*. PhD thesis, University of Munich, 2011. URL: <http://files.oliverfriedmann.de/theses/phd.pdf>.
- [FT87] Michael L. Fredman and Robert Endre Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM*, 34(3):596–615, 1987. doi:10.1145/28869.28874.
- [HK66] A. J. Hoffman and R. M. Karp. On nonterminating stochastic games. *Manage. Sci.*, 12(5):359–370, 1966. doi:10.1287/mnsc.12.5.359.
- [JL17] Marcin Jurdzinski and Ranko Lazic. Succinct progress measures for solving parity games. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS*, pages 1–9, 2017. doi:10.1109/LICS.2017.8005092.
- [JPZ08] Marcin Jurdzinski, Mike Paterson, and Uri Zwick. A deterministic subexponential algorithm for solving parity games. *SIAM J. Comput.*, 38(4):1519–1532, 2008. doi:10.1137/070686652.
- [Jur98] Marcin Jurdzinski. Deciding the winner in parity games is in $UP \cap co-UP$. *Inf. Process. Lett.*, 68(3):119–124, 1998. doi:10.1016/S0020-0190(98)00150-1.
- [Jur00] Marcin Jurdzinski. Small progress measures for solving parity games. In *17th Annual Symposium on Theoretical Aspects of Computer Science, STACS*, volume 1770 of *Lecture Notes in Computer Science*, pages 290–301, 2000. doi:10.1007/3-540-46541-3_24.
- [KKV01] Valerie King, Orna Kupferman, and Moshe Y. Vardi. On the complexity of parity word automata. In *Foundations of Software Science and Computation Structures, 4th International Conference, FOSSACS*, volume 2030 of *Lecture Notes in Computer Science*, pages 276–286. Springer, 2001. doi:10.1007/3-540-45315-6_18.
- [KV98] Orna Kupferman and Moshe Y. Vardi. Weak alternating automata and tree automata emptiness. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, STOC*, pages 224–233. ACM, 1998. doi:10.1145/276698.276748.
- [Leh18] Karoliina Lehtinen. A modal μ perspective on solving parity games in quasi-polynomial time. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS*, pages 639–648, 2018. doi:10.1145/3209108.3209115.
- [Lut08] Michael Luttenberger. Strategy iteration using non-deterministic strategies for solving parity games. *CoRR*, abs/0806.2923, 2008. URL: <http://arxiv.org/abs/0806.2923>, arXiv:0806.2923.
- [Mar75] Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975. doi:10.2307/1971035.
- [MRR18] Matthias Mnich, Heiko Röglin, and Clemens Rösner. New deterministic algorithms for solving parity games. *Discret. Optim.*, 30:73–95, 2018. doi:10.1016/J.DISOPT.2018.06.001.

- [Ohl21] Pierre Ohlmann. *Monotonic graphs for parity and mean-payoff games*. PhD thesis, Université Paris Cité, 2021. URL: <https://tel.archives-ouvertes.fr/tel-03771185>.
- [Par19] Pawel Parys. Parity games: Zielonka’s algorithm in quasi-polynomial time. In *44th International Symposium on Mathematical Foundations of Computer Science, MFCS*, volume 138 of *LIPICs*, pages 10:1–10:13, 2019. doi:10.4230/LIPICs.MFCS.2019.10.
- [Pur95] Anuj Puri. *Theory of hybrid systems and discrete event systems*. PhD thesis, EECS Department, University of California, Berkeley, 1995.
- [Sch08] Sven Schewe. An optimal strategy improvement algorithm for solving parity and payoff games. In *22nd International Workshop on Computer Science Logic, CSL*, volume 5213 of *Lecture Notes in Computer Science*, pages 369–384, 2008. doi:10.1007/978-3-540-87531-4_27.
- [Sch17] Sven Schewe. Solving parity games in big steps. *J. Comput. Syst. Sci.*, 84:243–262, 2017. doi:10.1016/J.JCSS.2016.10.002.
- [Tar72] Robert Endre Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1(2):146–160, 1972. doi:10.1137/0201010.
- [VJ00] Jens Vöge and Marcin Jurdzinski. A discrete strategy improvement algorithm for solving parity games. In *12th International Conference on Computer-Aided Verification, CAV*, volume 1855 of *Lecture Notes in Computer Science*, pages 202–215, 2000. doi:10.1007/10722167_18.
- [Zie98] Wieslaw Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theor. Comput. Sci.*, 200(1-2):135–183, 1998. doi:10.1016/S0304-3975(98)00009-7.

APPENDIX A. CONNECTION TO TROPICAL LINEAR PROGRAMMING

In this section, we describe the well-known connection between parity games and tropical linear programming. To this end, let us assume that G is bipartite with bipartition $V_0 \sqcup V_1$ without loss of generality. Consider the tropical (min-plus) semiring $(\mathbb{T}, \oplus, \odot)$, where the set $\mathbb{T} = \mathbb{R} \cup \{\infty\}$ is equipped with binary operations $a \oplus b = \min\{a, b\}$ and $a \odot b = a + b$. Let $n_1 = |V_1|$ and $y \in \mathbb{T}^{n_1}$ be a vector of variables indexed over the nodes in V_1 . A parity game can be formulated as the following system of tropical linear inequalities

$$y_u + (-n)^{\pi(u)} \geq \min_{vw \in E} \left\{ y_w - (-n)^{\pi(v)} \right\} \quad \forall uv \in \delta^+(V_1). \quad (\text{P})$$

This is precisely the reduction from parity games to mean payoff games [EM79, Pur95]. In particular, for each node $v \in V$, we assign a cost of $-(-n)^{\pi(v)}$ to every arc in $\delta^+(v)$. A trivial feasible solution to (P) is given by $\infty \cdot \mathbb{1}$. However, we are interested in feasible solutions of maximal finite support due to the following statement.

Theorem A.1 [AGG12]. *Let $y^* \in \mathbb{T}^{n_1}$ be a feasible solution to (P) of maximal finite support. Then, Even wins from $u \in V_1$ if and only if $y_u^* < \infty$.*