

Good-for-Games ω -Pushdown Automata

Lehtinen, Karoliina; Zimmermann, Martin

Published in:
Logical Methods in Computer Science

DOI (link to publication from Publisher):
[10.46298/LMCS-18\(1:3\)2022](https://doi.org/10.46298/LMCS-18(1:3)2022)

Creative Commons License
Unspecified

Publication date:
2022

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Lehtinen, K., & Zimmermann, M. (2022). Good-for-Games ω -Pushdown Automata. *Logical Methods in Computer Science*, 18(1). [https://doi.org/10.46298/LMCS-18\(1:3\)2022](https://doi.org/10.46298/LMCS-18(1:3)2022)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

GOOD-FOR-GAMES ω -PUSHDOWN AUTOMATA

KAROLIINA LEHTINEN^a AND MARTIN ZIMMERMANN^b

^a CNRS, Aix-Marseille Université and Université de Toulon, LIS, Marseille, France
e-mail address: karoliina.lehtinen@lis-lab.fr

^b University of Liverpool, Liverpool, United Kingdom
current affiliation: Aalborg University, Aalborg, Denmark
e-mail address: Martin.Zimmermann@liverpool.ac.uk

ABSTRACT. We introduce good-for-games ω -pushdown automata (ω -GFG-PDA). These are automata whose nondeterminism can be resolved based on the input processed so far. Good-for-gameness enables automata to be composed with games, trees, and other automata, applications which otherwise require deterministic automata.

Our main results are that ω -GFG-PDA are more expressive than deterministic ω -pushdown automata and that solving infinite games with winning conditions specified by ω -GFG-PDA is EXPTIME-complete. Thus, we have identified a new class of ω -contextfree winning conditions for which solving games is decidable. It follows that the universality problem for ω -GFG-PDA is in EXPTIME as well.

Moreover, we study closure properties of the class of languages recognized by ω -GFG-PDA and decidability of good-for-gameness of ω -pushdown automata and languages. Finally, we compare ω -GFG-PDA to ω -visibly PDA, study the resources necessary to resolve the nondeterminism in ω -GFG-PDA, and prove that the parity index hierarchy for ω -GFG-PDA is infinite.

1. INTRODUCTION

Good-for-gameness is the new determinism, and not just for solving games. Good-for-games automata also lend themselves to composition with other automata and trees. These problems have in common that they are traditionally addressed with deterministic automata, which, depending on the exact type used, may be less succinct or even less expressive than nondeterministic ones. Good-for-games automata overcome this restriction by allowing a limited form of nondeterminism that does not interfere with composition.

As an example, consider the setting of infinite-duration two-player zero-sum games of perfect information. In such a game, two players interact to produce a play, an infinite word over some alphabet. A winning condition specifies a partition of the set of plays indicating the winner of each play. Here, we are concerned with games whose winning condition is

Key words and phrases: Good-for-games, Pushdown Automata, Infinite Games.

A preliminary version of this manuscript was presented at LICS 2020 [LZ20]. This version extends it with an analysis of resource-bounded resolvers, of the parity index hierarchy as well as of the synthesis problem.

explicitly given by an automaton recognizing the winning plays for one player.¹ This setting arises, for example, when solving the LTL synthesis problem where the winning condition is specified by an LTL formula, which can be turned into an automaton.

The usual approach to solving a game with a winning condition given by an automaton \mathcal{A} is to obtain an equivalent deterministic automaton \mathcal{D} and then solve an arena-based game with an implicit winning condition given by the acceptance condition of \mathcal{D} . The arena simultaneously captures the interaction between the players, which results in a play, and constructs the run of \mathcal{D} on this play. The resulting arena-based game has the same winner as the original game with winning condition recognized by \mathcal{A} . For example, if the original winning condition is ω -regular, there is a deterministic parity automaton recognizing it, and the resulting game is a parity game, which can be effectively solved.

However, the correctness of this reduction crucially depends on the on-the-fly construction of the run of \mathcal{D} on the play. For nondeterministic automata, one might be tempted to let the nondeterministic choices be resolved by the player who wins if the automaton accepts. However, an accepting run cannot necessarily be constructed on-the-fly, even if one exists, as the resolution of nondeterministic choices might depend on the whole play rather than on the current finite play prefix. A simple example is a winning condition that allows the player who resolves the nondeterminism to win the original game while her opponent wins in the arena-based game by using her nondeterministic choices against her. This is the case in the parity automaton presented in Figure 1, which accepts all words, but in which nondeterminism must decide whether a word has finitely or infinitely many occurrences of a .




Figure 1: A nondeterministic parity automaton that recognises $\{a, b\}^\omega$. A run is accepting if it visits the black vertex only finitely often and a white vertex infinitely often.

Assume that the opponent has control over the number of a 's that appear during a play. This number does not affect the winner of the original game, as the automaton is universal. But if the initial nondeterminism is resolved by moving to q_1 , then the opponent wins by enforcing infinitely many a , while if the nondeterminism is resolved by moving to q_2 then he wins by enforcing finitely many a .

Good-for-games automata (also known as history-deterministic automata² [Col09]), introduced by Henzinger and Piterman [HP06], are nondeterministic (or even alternating [Col13, BL19]) automata whose nondeterminism can be resolved based only on the input processed so far. This property implies that the previously described procedure yields the correct winner, even if the automaton is not deterministic.

¹This should be contrasted with the classical setting of say parity games, where the winning condition is implicitly encoded by a colouring of the arena specifying the interaction between the players.

²The terms “history-deterministic” and “good-for-games” have often been used interchangeably, but recently it has come to light that history-determinism and compositionality with games do not coincide on all models [BL21]. We still prefer to use the term “good-for-games” for consistency with previous publications, although it is open whether these two notions coincide in the setting we consider here.

Since their introduction, Boker, Kupferman, Kuperberg and Skrzypczak have shown that ω -regular good-for-games automata are also suitable for composition with trees [BKKM13], which can be seen as the special case of one-player arena-based games, while Boker and Lehtinen have shown that good-for-games automata are suitable for automata composition in the following sense [BL19]: If an ω -regular good-for-games automaton \mathcal{B} recognizes the set of accepting runs of an alternating ω -regular automaton \mathcal{A} , then the composition of \mathcal{A} and \mathcal{B} is an ω -regular automaton that recognizes the same language as \mathcal{A} , but with the acceptance condition of \mathcal{B} . In other words, good-for-games automata, like deterministic automata, can be used both to simplify the winning conditions of games and the acceptance conditions of automata. Kuperberg and Skrzypczak showed that ω -regular good-for-games co-Büchi automata can be exponentially more succinct than deterministic ones while good-for-games Büchi automata are at most quadratically more succinct [KM15].

However, since deterministic parity automata, which are trivially good-for-games, express all ω -regular languages, succinctness is the most good-for-gameness can offer in the ω -regular setting. In contrast, deterministic automata models are in general less expressive, not just less succinct, than their nondeterministic counterparts. This is true, for example, for pushdown automata and for various types of quantitative automata. We argue that in such cases, it is worthwhile to investigate good-for-games automata as an alternative to deterministic ones, as they form a potentially larger class of winning conditions for which solving games is decidable. Indeed, the study of quantitative automata lead to the independent introduction of good-for-games automata by Colcombet. In particular, in the setting of regular cost functions, good-for-games cost automata are as expressive as nondeterministic ones, unlike deterministic ones [Col09].

So far the case of ω -pushdown automata has not been considered. Here, the increased expressiveness of nondeterministic automata comes at a heavy price: games with winning conditions given by nondeterministic ω -pushdown automata are undecidable [Fin01] while those with winning conditions given by deterministic ones are decidable [Wal01]. Hence, in this work, we introduce and study good-for-games ω -pushdown automata (ω -GFG-PDA) to push the frontier of decidability for games with ω -contextfree winning conditions.

Our contributions. Our first results concern expressiveness: In Section 3, we prove that ω -GFG-PDA are strictly more expressive than deterministic ω -pushdown automata (ω -DPDA), but not as expressive as nondeterministic ω -pushdown automata (ω -PDA). So, they do form a new class of ω -contextfree languages. This is in contrast to the ω -regular setting where, for each of the usual acceptance conditions, deterministic and GFG automata recognize exactly the same languages. Also, in Section 8, we prove that the parity index hierarchy is infinite for ω -GFG-PDA.

Second, in Section 4, we show that ω -GFG-PDA live up to their name: Determining the winner of a game with a winning condition specified by an ω -GFG-PDA is EXPTIME-complete, as for the special case of games with winning conditions specified by ω -DPDA [Wal01]. Furthermore, winning strategies for the player aiming to satisfy the specification are implementable by pushdown transducer, which can be computed in exponential time. The player violating the specification does not necessarily have a winning strategy that is implementable by a pushdown transducer. The decidability result has to be contrasted with the undecidability of games with a winning condition specified by an ω -PDA [Fin01]. As a corollary, the universality of ω -GFG-PDA is also in EXPTIME, while it is undecidable for ω -PDA. Table 1 sums up our results on decidability.

	Emptiness	Universality	Solving Games	
ω -DPDA	PTime	PTime	EXPTIME	[CG78, Wal01]
ω -GFG-PDA	PTime	EXPTIME*	EXPTIME	here
ω -VPA	PTime	EXPTIME	2EXPTIME	[AM04, LMS04]
ω -PDA	PTime	undecidable	undecidable	[CG77a, CG77b, Fin01]

Table 1: Summary of our results on decision problems (in light gray) and comparison to other classes of contextfree languages. All problems are complete for the respective complexity class unless marked with an asterisks.

	Intersection	Union	Complement	Set difference	Homomorphism
ω -DPDA	✗	✗	✓	✗	✗
ω -GFG-PDA	✗	✗	✗	✗	✗
ω -VPA	✓	✓	✓	✓	✗
ω -PDA	✗	✓	✗	✗	✓

Table 2: Summary of our results on closure properties (in light gray) and comparison to other classes of contextfree languages.

Third, in Section 5, we study the closure properties of ω -GFG-PDA, which are almost nonexistent, and, in Section 6, prove that both the problems of deciding whether a given ω -PDA is good-for games, and of deciding whether a given ω -PDA is language equivalent to an ω -GFG-PDA are undecidable. Table 2 sums up our results on closure properties.

Fourth, in Section 7, we study the resources necessary to resolve nondeterminism in ω -GFG-PDA. In general, pushdown transducers are not sufficient while ω -GFG-PDA with finite-state machines resolving the nondeterminism are always determinizable.

Fifth, in Section 9, we compare ω -GFG-PDA with visibly pushdown automata (ω -VPA) [AM04], a class of ω -PDA with robust closure properties and for which solving games is also decidable [LMS04]. We show that the classes of languages recognized by ω -GFG-PDA and ω -VPA are incomparable with respect to inclusion. See Figure 2 for an overview of our results on the relations between these classes.

2. PRELIMINARIES

An alphabet Σ is a finite nonempty set of letters. The set of finite words over Σ is denoted by Σ^* , the set of nonempty finite words over Σ by Σ^+ , and the set of infinite words over Σ by Σ^ω . The empty word is denoted by ε , the length of a finite word v is denoted by $|v|$, and the n -th letter of a finite or infinite word is denoted by $w(n)$ (starting with $n = 0$). An ω -language over Σ is a subset of Σ^ω .

For alphabets Σ_1, Σ_2 , we extend functions $f: \Sigma_1 \rightarrow \Sigma_2^*$ homomorphically to finite and infinite words over Σ_1 via

$$f(w) = f(w(0))f(w(1))f(w(2))\cdots.$$

For example, if $\Sigma_1 = \Sigma \times \Sigma'$, then $\pi_i(a_1, a_2) \mapsto a_i$ for $(a_1, a_2) \in \Sigma_1$ denotes the projection to the i -th component ($i \in \{1, 2\}$).

An ω -pushdown automaton (ω -PDA for short) $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, \Omega)$ consists of a finite set Q of states with the initial state $q_I \in Q$, an input alphabet Σ , a stack alphabet Γ , a transition relation Δ to be specified, and a coloring $\Omega: \Delta \rightarrow \mathbb{N}$. For notational convenience, we define $\Sigma_\varepsilon = \Sigma \cup \{\varepsilon\}$ and $\Gamma_\perp = \Gamma \cup \{\perp\}$, where $\perp \notin \Gamma$ is a designated stack bottom symbol. Then, the transition relation Δ is a subset of $Q \times \Gamma_\perp \times \Sigma_\varepsilon \times Q \times \Gamma_\perp^{\leq 2}$ that we require to neither write nor delete the stack bottom symbol from the stack: If $(q, \perp, a, q', \gamma) \in \Delta$, then $\gamma \in \perp \cdot (\Gamma \cup \{\varepsilon\})$, and if $(q, X, a, q', \gamma) \in \Delta$ for $X \in \Gamma$, then $\gamma \in \Gamma^{\leq 2}$. Given a transition $\tau = (q, X, a, q', \gamma)$ let $\ell(\tau) = a \in \Sigma_\varepsilon$. We say that τ is an $\ell(\tau)$ -transition and that τ is a Σ -transition, if $\ell(\tau) \in \Sigma$. For a finite or infinite sequence ρ over Δ , $\ell(\rho)$ is defined by applying ℓ homomorphically to every transition.

A stack content is a finite word in $\perp\Gamma^*$ (i.e., the top of the stack is at the end) and a configuration $c = (q, \gamma)$ of \mathcal{P} consists of a state $q \in Q$ and a stack content γ . The stack height of c is $\text{sh}(c) = |\gamma| - 1$. The initial configuration is (q_I, \perp) .

A transition $\tau = (q, X, a, q', \gamma') \in \Delta$ is enabled in a configuration c if $c = (q, \gamma X)$ for some $\gamma \in \Gamma_\perp^*$. In this case, we write $(q, \gamma X) \xrightarrow{\tau} (q', \gamma\gamma')$. A run of \mathcal{P} is an infinite sequence $\rho = c_0\tau_0c_1\tau_1c_2\tau_2\cdots$ of configurations and transitions with c_0 being the initial configuration and $c_n \xrightarrow{\tau_n} c_{n+1}$ for every n . Finite run prefixes are defined analogously and are required to end in a configuration. The infinite run ρ is a run of \mathcal{P} on $w \in \Sigma^\omega$, if $w = \ell(\tau_0\tau_1\tau_2\cdots)$ (this implies that ρ contains infinitely many Σ -transitions). We say that ρ is accepting if $\limsup_{n \rightarrow \infty} \Omega(\tau_n)$ is even, i.e., if the maximal color labeling infinitely many transitions is even. The language $L(\mathcal{P})$ recognized by \mathcal{P} contains all $w \in \Sigma^\omega$ such that \mathcal{P} has an accepting run on w .

Remark 2.1. Let $c_0\tau_0c_1\tau_1c_2\tau_2\cdots$ be a run of \mathcal{P} . Then, the sequence $c_0c_1c_2\cdots$ of configurations is uniquely determined by the sequence $\tau_0\tau_1\tau_2\cdots$ of transitions. Hence, whenever convenient, we treat a sequence of transitions as a run if it indeed induces one (not every such sequence does induce a run, e.g., if a transition τ_n is not enabled in c_n).

We say that an ω -PDA \mathcal{P} is deterministic if

- for every $q \in Q$, every $X \in \Gamma_\perp$, and every $a \in \Sigma_\varepsilon$, there is at most one transition of the form $(q, X, a, q', \gamma) \in \Delta$ for some q' and some γ , and




Figure 2: The classes of ω -languages recognized by the automata considered in this work. Languages separating these classes are shown as black dots.




Figure 3: The ω -PDA \mathcal{P} from Example 2.2. The self-loop at the white state q_4 has color 2 while all other transitions have color 1, and X represents an arbitrary stack symbol from the stack alphabet $\{A, B, N\}$ (not including \perp).

- for every $q \in Q$ and every $X \in \Gamma_\perp$, if there is a transition $(q, X, \varepsilon, q_1, \gamma_1) \in \Delta$ for some q_1 and some γ_1 , then there is no $a \in \Sigma$ such that there is a transition $(q, X, a, q_2, \gamma_2) \in \Delta$ for some q_2 and some γ_2 .

As expected, a deterministic ω -pushdown automaton (ω -DPDA) has at most one run on every ω -word.

The class of ω -languages recognized by ω -PDA is denoted by ω -CFL and the class of ω -languages recognized by ω -DPDA by ω -DCFL. Cohen and Gold showed that ω -DCFL is a strict subset of ω -CFL [CG78, Theorem 3.2].³

Example 2.2. The ω -PDA \mathcal{P} depicted in Figure 3 recognizes the ω -language

$$\{ac^n d^n \#^\omega \mid n \geq 1\} \cup \{bc^n d^{2n} \#^\omega \mid n \geq 1\}.$$

Note that while \mathcal{P} is nondeterministic, $L(\mathcal{P})$ is in ω -DCFL.

3. GOOD-FOR-GAMES PUSHDOWN AUTOMATA

Here, we introduce good-for-games ω -pushdown automata (ω -GFG-PDA for short), nondeterministic ω -pushdown automata whose nondeterminism can be resolved based on the run prefix constructed thus far and on the next input letter to be processed, but independently of the continuation of the input beyond the next letter.

As an example, consider the ω -PDA \mathcal{P} from Example 2.2. It is nondeterministic, but the nondeterminism in a configuration of the form $(q_1, \gamma N)$ can be resolved based on whether the first transition of the run processed an a or a b : in the former case, take the transition to state q_2 , in the latter case the transition to state q_3 . Afterwards, there are no nondeterministic choices to make and the resulting run is accepting whenever the input is in the language. This automaton is therefore good-for-games. The language of this automaton is clearly ω -DCFL; an automaton separating ω -DCFL and the class of languages recognised by ω -GFG-PDA is one of the main results of this section, presented below.

³Formally, Cohen and Gold considered automata with state-based Muller acceptance while we consider, for technical convenience, automata with transition-based parity acceptance. However, using *latest appearance records* (see, e.g., [GTW02]) shows that both definitions are equivalent.

Formally, we say that an ω -PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, \Omega)$ is good-for-games, if there is a (nondeterminism) resolver for \mathcal{P} , a function $r: \Delta^* \times \Sigma \rightarrow \Delta$ such that for every $w \in L(\mathcal{P})$ the sequence $\tau_0\tau_1\tau_2 \cdots \in \Delta^\omega$ defined by

$$\tau_n = r(\tau_0 \cdots \tau_{n-1}, w(|\ell(\tau_0 \cdots \tau_{n-1})|))$$

induces an accepting run of \mathcal{P} on w . Note that the prefix processed so far can be recovered from r 's input, i.e., it is $\ell(\tau_0 \cdots \tau_{n-1})$. However, the converse is not true due to the existence of ε -transitions. This is the reason the run prefix and not the input prefix is the input for the resolver.

Remark 3.1. Let \mathcal{P} be an ω -GFG-PDA over Σ such that every finite word $v \in \Sigma^*$ is a prefix of some word in $L(\mathcal{P})$, and let r be a resolver for \mathcal{P} .

Then, r induces a run on every input $w \in \Sigma^\omega$, even if w is not accepted by \mathcal{P} . Indeed, if there is no run on w , because either there is no enabled transition that processes the next input letter or because it ends in an infinite tail of ε -transitions, then some word in $L(\mathcal{P})$ has no accepting run induced by r , which is a contradiction.

Each deterministic automaton is trivially good-for-games. We denote the class of ω -languages recognized by ω -GFG-PDA by ω -GFG-CFL. By definition, the three classes of languages we consider form a hierarchy.

Proposition 3.2. ω -DCFL \subseteq ω -GFG-CFL \subseteq ω -CFL.

We show that both inclusions are strict. In particular, ω -GFG-PDA are more expressive than ω -DPDA.

Let $I = \{0, +, -\}$ and define the energy level $EL(v) \in \mathbb{Z}$ of finite words v over I inductively as $EL(\varepsilon) = 0$, and $EL(v0) = EL(v)$, $EL(v+) = EL(v) + 1$, as well as $EL(v-) = EL(v) - 1$. We say that a word $w \in I^\omega$ is *safe* if $EL(w(0) \cdots w(n)) \geq 0$ for every $n \geq 0$.

Remark 3.3. Let $w \in I^\omega$.

- (1) w has a safe suffix if and only if there is an $s \in \mathbb{N}$ such that $EL(w(0) \cdots w(n)) \geq -s$ for all n .
- (2) If w is safe then there is an $n > 0$ such that $w(n)w(n+1)w(n+2) \cdots$ is safe as well.

We fix $\Sigma = I \times I$ and define L_{ss} to be the language containing all $w \in \Sigma^\omega$ such that $\pi_i(w)$ has a safe suffix, for some $i \in \{1, 2\}$.

Lemma 3.4. $L_{ss} \in \omega$ -GFG-CFL.

Proof. Consider the ω -PDA \mathcal{P} in Figure 4 with two states 1 and 2 signifying whether a potential safe suffix is being tracked in the first or second component of the input, and a single stack symbol N used to track the energy level of such a suffix. The initial state is arbitrary; we fix it to be 1.

The automaton can, at any moment, nondeterministically change its state from 1 to 2 and vice versa without changing the stack content (while processing an input letter that is just ignored). When not changing its state, say while staying in state i , \mathcal{P} deterministically processes the next input letter $\binom{a_1}{a_2}$. If $a_i = 0$ then the stack is left unchanged and if $a_i = +$ then an N is pushed onto the stack. If $a_i = -$ and the stack is nonempty, then the topmost N is popped from the stack. The stack is left unchanged if $a_i = -$ and the stack is empty. Note that the only nondeterministic choice is to change the state, i.e., if the state is not changed then the automaton has exactly one transition that can process the next input letter. Furthermore, \mathcal{P} has no ε -transitions, i.e., each transition processes an input letter.




Figure 4: The ω -PDA \mathcal{P} recognizing L_{ss} . Here, $*$ is an arbitrary letter in $\{0, +, -\}$ and X an arbitrary stack symbol from the stack alphabet $\{N\}$ or \perp . Transitions with light gray background have color 1, all others have color 0.

A transition has color 1 if it either changes the state, or if it processes a $-$ in the i -th component while in state i with an empty stack. All other transitions have color 0. Hence, a run is accepting if and only if it eventually stays in some state i and then does not process a $-$ in component i while the stack is empty.

First, we show that $L_{ss} = L(\mathcal{P})$, then that \mathcal{P} is good-for-games. Let $w \in L_{ss}$, i.e., there is an $i \in \{1, 2\}$ such that $\pi_i(w) = a(0)a(1)a(2) \cdots$ has a safe suffix, say $a(n)a(n+1)a(n+2) \cdots$. Due to Remark 3.3(2), we assume w.l.o.g. $n > 0$. Consider the unique run $\rho = c_0\tau_0c_1\tau_1c_2\tau_2 \cdots$ of \mathcal{P} on w that immediately switches to state i (if necessary) and otherwise always executes the unique enabled transition that processes the next input letter without changing state. We have $\ell(\tau_0 \cdots \tau_{j-1}) = w(0) \cdots w(j-1)$ due to the absence of ε -transitions.

An induction shows the invariant $\text{sh}(c_{n+j}) = \text{sh}(c_n) + EL(a(n) \cdots a(n+j-1))$ for every $j \geq 0$. Here we use the assumption $n > 0$, which ensures that no transitions changing the state are used to process the safe suffix. Hence, the state is equal to i at all but possibly the first configuration and the invariant implies that no $-$ is processed in component i by a transition τ_{n+j} while the stack is empty. This implies that the run is accepting, i.e., $w \in L(\mathcal{P})$.

Conversely, let $w \in L(\mathcal{P})$. Then, there is an accepting run $c_0\tau_0c_1\tau_1c_2\tau_2 \cdots$ of \mathcal{P} on w . Again, as \mathcal{P} has no ε -transitions, $\ell(\tau_0 \cdots \tau_{j-1}) = w(0) \cdots w(j-1)$. Now, let τ_{n-1} be the last transition with color 1 (pick $\tau_{n-1} = \tau_0$ if there is no such transition) and let $c_n = (i, \perp N^s)$. We define $a(n)a(n+1)a(n+2) \cdots$ to be the suffix of $\pi_i(w) = a(0)a(1)a(2) \cdots$ starting at position n .

By the choice of n , after τ_{n-1} no transition changes the state (it is always equal to i) and there is no $j \geq 0$ such that $a(n+j) = -$ and $\text{sh}(c_{n+j}) = 0$. Thus, an induction shows that $EL(a(n) \cdots a(n+j)) = \text{sh}(c_{n+j}) - s$. Hence, the energy level of the prefixes of $a(n)a(n+1)a(n+2) \cdots$ is bounded from below by $-s$. Thus, Remark 3.3(1) implies that $a(n)a(n+1)a(n+2) \cdots$ has a safe suffix, i.e., $w \in L_{ss}$.

It remains to show that \mathcal{P} is good-for-games.⁴ Intuitively, we construct a resolver r that always chooses to track the input component that has the longest suffix that can still be extended to an infinite safe word (preferring the first component in case of a tie). If one of the components has a safe suffix, eventually the resolver will choose to track it. The resulting run will be accepting.

⁴Although the language we use here is different, this argument is similar to the one used by Kuperberg and Skrzypczak [KM15] to show that good-for-games co-Büchi automata are exponentially more succinct than deterministic ones.

Fix $v = \binom{a_1(0)}{a_2(0)} \cdots \binom{a_1(n)}{a_2(n)} \in \Sigma^+$ and let S_i contain those $j \leq n$ such that $a_i(j) \cdots a_i(n)$ is safe (which is defined as expected). Further, let

$$i(v) = \begin{cases} 1 & \text{if } \min S_1 \leq \min S_2, \\ 2 & \text{otherwise,} \end{cases}$$

where we use $\min \emptyset = n + 1$. We define $r(\tau_0 \cdots \tau_{n-1}, a)$ inductively to always ensure that its target state is equal to $i(\ell(\tau_0 \cdots \tau_{n-1})a)$ and, if this does not require a state change, then the unique transition processing the $i(\ell(\tau_0 \cdots \tau_{n-1})a)$ -th component of a is returned.

Now, let $w \in L_{ss}$, i.e., there is an $i \in \{1, 2\}$ such that $\pi_i(w)$ has a safe suffix. Then, r produces a run that tracks the safe suffix that starts as early as possible (again favoring the first component in case of a tie). As in the argument above one can show that this run is accepting, as it switches states only finitely often and does not process a $-$ while the stack is empty while tracking the safe suffix. Hence, r has the desired properties and \mathcal{P} is good-for-games. \square

After having shown that L_{ss} is in ω -GFG-CFL, we show that it is not in ω -DCFL, thereby separating ω -DCFL and ω -GFG-CFL.

Lemma 3.5. $L_{ss} \notin \omega$ -DCFL.

Proof. We assume towards a contradiction that there is an ω -DPDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, \Omega)$ recognizing L_{ss} . Define $x_1 = \binom{+}{0} \binom{+}{-}$ and $x_2 = \binom{0}{+} \binom{-}{+}$, i.e., in x_i the energy level in component i is increased by two while it is decreased by one in the other component.

Define

$$w_{\overline{ss}} = x_1 (x_2)^3 (x_1)^7 (x_2)^{15} (x_1)^{31} (x_2)^{63} \cdots$$

An induction shows

$$EL(\pi_1(x_1(x_2)^3 \cdots (x_2)^{2^{2j-1}})) = -j$$

for every $j > 1$ and

$$EL(\pi_2(x_1(x_2)^3 \cdots (x_1)^{2^{2j-1}-1})) = -j$$

for every $j > 0$. Hence, due to Remark 3.3(1), $w_{\overline{ss}} \notin L_{ss}$.

Due to Remark 3.1, \mathcal{P} has a run $\rho = c_0 \tau_0 c_1 \tau_1 c_2 \tau_2 \cdots$ on $w_{\overline{ss}}$, which is rejecting. A step of ρ is a position n such that $\text{sh}(c_n) \leq \text{sh}(c_{n+j})$ for all $j \geq 0$. Every infinite run has infinitely many steps. Hence, we can find two steps $s < s'$ satisfying the following properties:

- (1) There is a state $q \in Q$ and a stack symbol $X \in \Gamma_\perp$ such that $c_s = (q, \gamma X)$ and $c_{s'} = (q, \gamma' X)$ for some γ, γ' , i.e., both configurations have the same state and topmost stack symbol.
- (2) The maximal color labeling the sequence $\tau_s \cdots \tau_{s'-1}$ of transitions leading from c_s to $c_{s'}$ is odd.
- (3) The sequence $\tau_s \cdots \tau_{s'-1}$ processes an infix v of w with $EL(\pi_i(v)) > 0$, for some $i \in \{1, 2\}$. This can be achieved, as every infix of length at least 3 of a word built by concatenating copies of the x_i has a strictly positive energy level in one component (note that the infix may start or end within an x_i).

Consider the sequence $\tau_0 \cdots \tau_{s-1} (\tau_s \cdots \tau_{s'-1})^\omega$ of transitions. Due to the first property, it induces a run ρ' of \mathcal{P} , which is rejecting due to the second property. Finally, due to the third property, ρ' processes a word with suffix v^ω . Such a word has a safe suffix in component i , as $EL(\pi_i(v)) > 0$. Hence, we have constructed a word in L_{ss} such that the

unique (as \mathcal{P} is deterministic) run of \mathcal{P} on w is rejecting, obtaining the desired contradiction to $L(\mathcal{P}) = L_{ss}$. \square

Our main result of this section is now a direct consequence of the previous two lemmata: ω -GFG-PDA are more expressive than ω -DPDA.

Theorem 3.6. ω -DCFL \subsetneq ω -GFG-CFL.

The next obvious question is whether every (nondeterministic) ω -contextfree language is good-for-games. Not unexpectedly, this is not the case. The intuitive reason is that good-for-games automata allow to resolve nondeterminism based on the history of a run, but still cannot resolve nondeterminism based on the continuation of the input. Considering a language that requires nondeterministic choices about the continuation of the input yields the desired separation. To this end, we adapt the classical proof that

$$\{a^n b^n \mid n \geq 1\} \cup \{a^n b^{2n} \mid n \geq 1\}$$

is not recognizable by a DPDA over finite words to our setting. Recall that this proof proceeds by contradiction as follows: Consider a DPDA \mathcal{P} recognizing this language and inputs $w_0 = a^n b^n$ and $w_1 = a^n b^{2n}$, both in the language. The run of \mathcal{P} on w_0 is a prefix of the run ρ_1 of \mathcal{P} on w_1 and ends an accepting state. Furthermore, no other prefix of ρ_1 can be accepting.

So, one can modify \mathcal{P} so that it behaves like \mathcal{P} until an accepting state is reached for the first time (the run thus far has processed an input $a^n b^n$ for some n up to that point due to the prefix property above). From that point onward, the modified automaton processes c 's instead of b 's (but behaves like \mathcal{P} otherwise) until it visits an accepting state again, i.e., instead of processing another n b 's it processes n c 's. Hence, the modified automaton accepts the non-contextfree language $\{a^n b^n c^n \mid n \geq 1\}$, which yields the desired contradiction.

Note that the prefix property of runs underlying the argument above is also satisfied by good-for-games automata. For technical convenience, we use the language

$$L = \{(a\#)^n (b\#)^n \#^\omega \mid n \geq 1\} \cup \{(a\#)^n (b\#)^{2n} \#^\omega \mid n \geq 1\},$$

i.e., we add the dummy symbols $\#$ which obfuscate the next (proper) letter to be processed from the resolver.

Theorem 3.7. ω -GFG-CFL \subsetneq ω -CFL.

Proof. The language L is in ω -CFL, as an ω -PDA can process a prefix $(a\#)^n$ by storing n in unary on the stack and then nondeterministically guess and verify whether the remaining suffix is $(b\#)^n \#^\omega$ (by popping a symbol from the stack for every b) or whether it is $(b\#)^{2n} \#^\omega$ (by popping a symbol from the stack for every other b), similarly to the automaton from Example 2.2

We claim that L is not in ω -GFG-CFL. We assume towards a contradiction that there is an ω -GFG-PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, \Omega)$ with $L(\mathcal{P}) = L$, say with resolver r . In what follows, we will reach a contradiction by constructing an ω -PDA that recognizes the language

$$L_{abc} = \{(a\#)^n (b\#)^n (c\#)^n \#^\omega \mid n \geq 1\},$$

which is not in ω -CFL. This follows from $\{a^n b^n c^n \mid n \geq 1\}$ not being contextfree and from the closure properties of ω -CFL [CG77b] and of the contextfree languages [HMu07].

First, we note that the language

$$C = \{\gamma q \in \perp \Gamma^* Q \mid \mathcal{P} \text{ accepts } \#^\omega \text{ when starting in } (q, \gamma)\}$$

is regular. This can be shown by first noting that

$$C_0 = \{\gamma Xq \in \perp \Gamma^* Q \mid \mathcal{P} \text{ accepts } \#^\omega \text{ when starting in } (q, \gamma X) \text{ with a run that only visits configurations of stack height greater or equal to } \text{sh}(q, \gamma X)\}$$

is a finite union of languages $\Gamma^* Xq \cap \perp \Gamma^* q$ for some $X \in \Gamma_\perp$ and some $q \in Q$, and therefore regular. Now, C is equal to

$$\{\gamma q \in \perp \Gamma^* Q \mid \text{there is a run infix } \rho \text{ with } \ell(\rho) \in \#^* \text{ leading from } (q, \gamma) \text{ to } C_0\}.$$

An application of standard saturation techniques [Bü64]⁵ (applied to the restriction of \mathcal{P} to transitions labeled by $\#$ or ε) shows that the latter set is regular, as the target set C_0 is regular.

Using a deterministic finite automaton $\mathcal{A} = (Q', \Gamma_\perp \cup Q, q'_I, \delta, F)$ recognizing C we construct an ω -PDA \mathcal{P}' as follows (also, see Figure 5): We extend the stack alphabet Γ of \mathcal{P} to $\Gamma \times Q'$ and define the transition relation of \mathcal{P}' so that it simulates a run of \mathcal{P} and keeps track of the state of \mathcal{A} reached by processing the stack content, i.e., if \mathcal{P}' reaches a stack content

$$\perp(X_1, q_1) \cdots (X_s, q_s)$$

then we have $q_j = \delta^*(q'_I, \perp X_0 \cdots X_j)$ for every $0 \leq j \leq s$.⁶

Additionally, the states of \mathcal{P}' have a Boolean flag that is changed if \mathcal{P}' leaves a configuration of the form (q, γ) with $\gamma q \in C$ for the first time. As the stack symbols of \mathcal{P}' encode the run of \mathcal{A} on the stack content, this can be checked easily. From there on, \mathcal{P}' continues to simulate a run of \mathcal{P} , but now every transition labeled by a b in \mathcal{P} is labeled by a c .

Finally, we define the acceptance condition of \mathcal{P}' such that it only accepts if it switches the flag, afterwards continues the simulation of a run of \mathcal{P} that is accepting, and processes at least one c after the switch. Note that this requires adding a second Boolean flag to the states to check whether a c has been processed. We claim that \mathcal{P}' recognizes the language L_{abc} .

To this end, let

$$w = (a\#)^n (b\#)^n (c\#)^n \#^\omega \in L_{abc}$$

and define

$$w_1 = (a\#)^n (b\#)^n \#^\omega \quad \text{and} \quad w_2 = (a\#)^n (b\#)^{2n} \#^\omega,$$

which are both in L . Thus, let ρ_1 and ρ_2 be the accepting runs of \mathcal{P} on w_1 and w_2 induced by the resolver r . A prefix ρ'_1 of ρ_1 processing $w'_1 = (a\#)^n (b\#)^{n-1} b$ is also a prefix of ρ_2 processing w'_1 (note that we have removed the last $\#$, as the resolver inducing the runs has access to the next letter to be processed). As ρ_1 is an accepting run of \mathcal{P} on $w_1 = w'_1 \#^\omega$, the last configuration of ρ'_1 is in C .

Hence, \mathcal{P}' can simulate the run prefix ρ'_1 processing w'_1 , switch the first flag and then continue to simulate the suffix of ρ'_2 obtained by removing ρ'_1 , which processes $\#(c\#)^n \#^\omega$. This run is accepting, i.e., we have $w = w'_1 \#(c\#)^n \#^\omega \in L(\mathcal{P}')$.

Now, let $w \in L(\mathcal{P}')$, i.e., there is an accepting run ρ' of \mathcal{P}' on w . By construction of \mathcal{P}' , we can split ρ' into a finite prefix ρ'_p before the first flag is switched and the corresponding infinite suffix ρ'_s starting with the switch. Again, by construction, ρ'_p is the simulation of a

⁵Also, see the survey of Carayol and Hague [CH14] for more details.

⁶Note that the simplest way to implement this is to replace each transition that swaps the topmost stack symbol from X to X' by two transitions, the first popping X from the stack, and the second pushing X' onto the stack (using a fresh state reached between the new transitions).




Figure 5: The construction of \mathcal{P}' (bottom, for the sake of readability, we only sketch the letter processed by a transition, not the stack operations, and unreachable states are in gray) from \mathcal{P} (top). We assume that C contains all configurations with a white state. A run proceeds from the left copy to the middle copy if the current configuration of \mathcal{P} is in C . In the middle and right copy, every b -label on a transition is replaced by a c . All transitions in the left and middle component are rejecting, while the transitions of the right copy have the same colors as in \mathcal{P} .

run ρ_p of \mathcal{P} that processes the same input and ends in a configuration in C , and no prefix of ρ_p ends in C . Hence, we can conclude that both ρ_p and ρ'_p process $(a\#)^n(b\#)^{n-1}b$ for some $n > 0$, as these are the minimal words leading to a configuration from which $\#^\omega$ can be accepted.

Now, consider the suffix ρ'_s , which processes at least one c . It also simulates a run suffix ρ_s of \mathcal{P} and $\ell(\rho'_s)$ is obtained from $\ell(\rho_s)$ by replacing each c by a b . Furthermore, ρ_s starts in the last configuration of ρ_p and satisfies the acceptance condition, as ρ'_s satisfies the acceptance condition. Hence, ρ_s processes $\#(b\#)^n\#^\omega$, as the concatenation of ρ_p and ρ_s is an accepting run. Altogether, ρ'_p processes $(a\#)^n(b\#)^{n-1}b$ and ρ'_s processes $\#(c\#)^n\#^\omega$, i.e., $w = (a\#)^n(b\#)^n(c\#)^n\#^\omega$, which is in L_{abc} .

We conclude that $L(\mathcal{P}') = L_{abc}$, which contradicts $L_{abc} \notin \omega\text{-CFL}$. Therefore, L separates $\omega\text{-GFG-CFL}$ and $\omega\text{-CFL}$. \square

The proof of Theorem 9.1 is based on another language separating $\omega\text{-GFG-CFL}$ and $\omega\text{-CFL}$. A third such language is the following, based on palindromes. Let $h: \{0, 1, \#\}^* \rightarrow \{0, 1\}^*$ be the homomorphism induced by $h(0) = 0$, $h(1) = 1$, and $h(\#) = \varepsilon$. Define

$$P = \{v\#^\omega \mid h(v) = xx^R \text{ for some } x \in \{0, 1\}^*\},$$

where x^R denotes the reversal of x . The proof that P indeed separates $\omega\text{-CFL}$ and $\omega\text{-GFG-CFL}$ can be found in the appendix.

4. GOOD-FOR-GAMES PUSHDOWN AUTOMATA ARE INDEED GOOD FOR GAMES

In this section, we show that the winner of infinite-duration games with ω -GFG-CFL winning conditions can be effectively determined. This result is best phrased in terms of Gale-Stewart games, abstract games without an arena [GS53], as we are interested in the influence of the winning condition on the decidability of solving games.⁷

Formally, a Gale-Stewart game $\mathcal{G}(L)$ is given by an ω -language $L \subseteq (\Sigma_1 \times \Sigma_2)^\omega$. It is played between Player 1 and Player 2 in rounds $n = 0, 1, 2, \dots$. In each round, first Player 1 picks a letter $a_1(n) \in \Sigma_1$, then Player 2 picks a letter $a_2(n) \in \Sigma_2$. After ω rounds, the players have constructed an outcome

$$w = \begin{pmatrix} a_1(0) \\ a_2(0) \end{pmatrix} \begin{pmatrix} a_1(1) \\ a_2(1) \end{pmatrix} \begin{pmatrix} a_1(2) \\ a_2(2) \end{pmatrix} \cdots$$

which is winning for Player 2 if it is in L . A strategy for Player 2 in $\mathcal{G}(L)$ is a mapping $\sigma: \Sigma_1^+ \rightarrow \Sigma_2$. The outcome w is consistent with σ , if $a_2(n) = \sigma(a_1(0) \cdots a_1(n))$ for all n . A strategy σ for Player 2 is winning if every outcome that is consistent with σ is in L . Player 2 wins $\mathcal{G}(L)$ if she has a winning strategy for $\mathcal{G}(L)$.

Proposition 4.1 [Fin01, Wal01].

- (1) *The following problem is undecidable: Given an ω -PDA \mathcal{P} , does Player 2 win $\mathcal{G}(L(\mathcal{P}))$?*
- (2) *The following problem is EXPTIME-complete: Given an ω -DPDA \mathcal{P} , does Player 2 win $\mathcal{G}(L(\mathcal{P}))$?*

Walukiewicz's decidability result [Wal01] is formulated for parity games on configuration graphs of pushdown automata. However, a Gale-Stewart game with ω -DCFL winning condition can be reduced in polynomial time to a parity game on a configuration graph of a pushdown machine, and vice versa. This construction crucially depends on the determinism of the automaton recognizing the winning condition, as witnessed by the undecidability result for winning conditions recognized by (possibly nondeterministic) ω -PDA.

Our main result shows that decidability extends to games given by ω -GFG-PDA, i.e., not all types of nondeterminism lead to undecidability.

Theorem 4.2. *The following problem is EXPTIME-complete: Given an ω -GFG-PDA \mathcal{P} , does Player 2 win $\mathcal{G}(L(\mathcal{P}))$?*

Proof. Given $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, \Omega)$ with $\Sigma = \Sigma_1 \times \Sigma_2$ we construct an ω -DPDA \mathcal{P}_d such that Player 2 wins $\mathcal{G}(L(\mathcal{P}))$ if and only if she wins $\mathcal{G}(L(\mathcal{P}_d))$. This yields EXPTIME membership, as determining the winner of Gale-Stewart games with ω -DCFL winning conditions is in EXPTIME (see Proposition 4.1(2)) and the size of \mathcal{P}_d is polynomial in the size of \mathcal{P} . The matching lower bound is immediate, as determining the winner of games with ω -DCFL winning conditions is already EXPTIME-hard (again, see Proposition 4.1(2)).

Intuitively, we construct an ω -DPDA \mathcal{P}_d that processes simultaneously both an input of \mathcal{P} and a run of \mathcal{P} , and checks whether the run is indeed an accepting run of \mathcal{P} on the input. In the corresponding Gale-Stewart game with winning condition $L(\mathcal{P}_d)$, Player 2 has to both choose a letter in Σ_2 and transitions of \mathcal{P} . In other words, we have moved the nondeterminism of \mathcal{P} into Player 2's moves. Then a winning strategy for Player 2 in $\mathcal{G}(L(\mathcal{P}))$ can be combined with a resolver to form a winning strategy in $\mathcal{G}(L(\mathcal{P}_d))$ and vice versa.

⁷Games in finite arenas can easily be encoded as Gale-Stewart games.

Formally, we construct \mathcal{P}_d such that it recognizes all ω -words $m_0 m_1 m_2 \cdots$ over $\Sigma_1 \times (\Sigma_2 \cup \Delta)$ where each *block* m_j is of the form

$$\begin{pmatrix} a_1(j) \\ a_2(j) \end{pmatrix} \begin{pmatrix} b_{j,0} \\ \tau_{j,0} \end{pmatrix} \cdots \begin{pmatrix} b_{j,n_j-1} \\ \tau_{j,n_j-1} \end{pmatrix} \begin{pmatrix} b_{j,n_j} \\ \tau_{j,n_j} \end{pmatrix}$$

for some $n_j \geq 0$ satisfying the following conditions:

- (1) The transitions $\tau_{j,0}, \dots, \tau_{j,n_j-1}$ are ε -transitions, and τ_{j,n_j} is an $\begin{pmatrix} a_1(j) \\ a_2(j) \end{pmatrix}$ -transition.
- (2) The sequence $\tau_{0,0} \cdots \tau_{0,n_0} \tau_{1,0} \cdots \tau_{1,n_1} \tau_{2,0} \cdots \tau_{2,n_2} \cdots$ of transitions induces an accepting run of \mathcal{P} on the ω -word $\begin{pmatrix} a_1(0) \\ a_2(0) \end{pmatrix} \begin{pmatrix} a_1(1) \\ a_2(1) \end{pmatrix} \begin{pmatrix} a_1(2) \\ a_2(2) \end{pmatrix} \cdots$. Note that all the $b_{j,j'}$ picked by Player 1 are ignored while Player 2 constructs the run, only the letters $a_1(j)$ are relevant.

If w is of that form then the decomposition into blocks is unique.

An ω -DPDA \mathcal{P}_d recognizing this language can easily be constructed in polynomial time from \mathcal{P} . To this end, \mathcal{P}_d deterministically simulates the transitions given in a block on the letter from $\Sigma_1 \times \Sigma_2$ at the beginning of the block. If a transition is not applicable, then the run terminates and is therefore rejecting. Some standard constructions are necessary to ensure that the input has the right format; in particular, we need to adapt the coloring to rule out that from some point onwards only ε -transitions appear in the input. It remains to show that Player 2 wins $\mathcal{G}(L(\mathcal{P}))$ if and only if she wins $\mathcal{G}(L(\mathcal{P}_d))$.

First, we construct a mapping $\sigma \mapsto \sigma_d$ turning a winning strategy σ for Player 2 in $\mathcal{G}(L(\mathcal{P}))$ into a winning strategy σ_d for Player 2 in $\mathcal{G}(L(\mathcal{P}_d))$. To this end, we fix a resolver $r: \Delta^* \times \Sigma \rightarrow \Delta$ for \mathcal{P} .

Intuitively, the strategy σ_d alternates between simulating a move of σ and then uses r to construct a sequence of transitions that processes the letter determined by the move. This sequence starts with a finite number of ε -transitions followed by one transition processing the letter.

More formally, define σ_d inductively starting with $\sigma_d(a) = \sigma(a)$ for $a \in \Sigma_1$. Now, let $v = a_1(0) \cdots a_1(n) \in \Sigma_1^*$ with $n > 0$ be an input such that

$$a_2(j) = \sigma_d(a_1(0) \cdots a_1(j))$$

is already defined for every $j < n$. To define $\sigma_d(v)$ we consider two cases.

If $a_2(n-1) \in \Sigma_2 \cup \Delta$ is a non- ε -transition, then we define $\sigma_d(v) = \sigma(v')$, where v' is obtained from v by removing the letters at positions j with $a_2(j) \in \Delta$. This simulates the next move of σ , as the transition $a_2(n-1)$ has processed the last letter. In the other case (i.e., if $a_2(n-1)$ is either a letter in Σ_2 or an ε -transition) define $\sigma_d(v) = r(\rho, \begin{pmatrix} a_1(j') \\ a_2(j') \end{pmatrix})$ where $\rho \in \Delta^*$ is obtained from $a_2(0) \cdots a_2(n-1)$ by removing the letters at positions j with $a_2(j) \in \Sigma_2$ and where $j' < n$ is maximal with $a_2(j') \in \Sigma_2$. This move continues the construction of a run infix that processes the last letter from $\Sigma_1 \times \Sigma_2$, which appears at position j' .

Now, let $w' \in (\Sigma_1 \times (\Sigma_2 \cup \Delta))^\omega$ be consistent with σ_d . An induction shows that w' is a sequence of blocks that encodes an outcome w over $\Sigma_1 \times \Sigma_2$ that is consistent with σ , and a run ρ of \mathcal{P} on w induced by r . As σ is a winning strategy, w is in $L(\mathcal{P})$, which implies that ρ is accepting. Hence, w' is in $L(\mathcal{P}_d)$, i.e., σ_d is indeed winning for Player 2 in $\mathcal{G}(L(\mathcal{P}_d))$.

Conversely, we construct a mapping $\sigma \mapsto \sigma_{-d}$ turning a winning strategy σ for Player 2 in $\mathcal{G}(L(\mathcal{P}_d))$ into a winning strategy σ_{-d} for Player 2 in $\mathcal{G}(L(\mathcal{P}))$. Intuitively, to define σ_{-d} , we simulate a play in $\mathcal{G}(L(\mathcal{P}))$ by a play in $\mathcal{G}(L(\mathcal{P}_d))$, and copy the choice of letters made by σ while ignoring the moves building the run of \mathcal{P} .

Fix some $\sqcup \in \Sigma_1$, which we use as dummy input. We inductively define for every input $v \in \Sigma_1^+$ for σ_{-d} an input $v_d \in \Sigma_1^+$ for σ and then define $\sigma_{-d}(v) = \sigma(v_d)$. We begin by defining $a_d = a$ for every $a \in \Sigma_1$. Now, assume we have defined v_d for some v . Let n be minimal such that $\sigma(v_d \sqcup^n) \in \Sigma_2 \cup \Delta$ is a non- ε -transition. Then, we define $(va)_d = v_d \sqcup^n a$. Intuitively, we extend v_d by irrelevant inputs until σ completes the run infix processing the last letter.

Let $w \in (\Sigma_1 \times \Sigma_2)^\omega$ be consistent with σ_{-d} . An induction shows that there is an $w' \in (\Sigma_1 \times (\Sigma_2 \cup \Delta))^\omega$ that is consistent with σ that encodes w and an accepting run of \mathcal{P} on w . Hence, $w \in L(\mathcal{P})$, i.e., σ_{-d} is indeed winning for Player 2 in $\mathcal{G}(L(\mathcal{P}))$. \square

Recall that a game is determined if either of the players has a winning strategy. Translating a winning strategy for Player 1 (defined as expected) for $\mathcal{G}(L(\mathcal{P}_d))$ into a winning strategy for him in $\mathcal{G}(L(\mathcal{P}))$ implies that $\mathcal{G}(L(\mathcal{P}))$ is determined, as $\mathcal{G}(L(\mathcal{P}_d))$ is determined. Such a transformation can be obtained along the lines of the transformations presented above for Player 2 winning strategies.

As universality of $L \subseteq \Sigma^\omega$ is equivalent to Player 2 winning $\mathcal{G}(\{(\frac{w}{\#^\omega}) \mid w \in L\})$ and as $\{(\frac{w}{\#^\omega}) \mid w \in L\}$ is in ω -GFG-CFL if L is in ω -GFG-CFL, we obtain the following corollary of our main theorem.

Corollary 4.3. *The following problem is in EXPTIME: Given an ω -GFG-PDA \mathcal{P} , is $L(\mathcal{P})$ universal?*

This contrasts with the universality problem for ω -PDA, which is undecidable. Emptiness of ω -GFG-PDA is also decidable, as it is decidable for ω -PDA, while equivalence of ω -GFG-PDA is undecidable, as it is already undecidable for ω -DPDA with Büchi or co-Büchi acceptance conditions [BGHH17].

Also, let us mention that one can apply the reduction presented in the proof of Theorem 4.2 also if \mathcal{P} is not known to be good-for-games. If Player 2 wins $\mathcal{G}(L(\mathcal{P}_d))$, then she wins $\mathcal{G}(L(\mathcal{P}))$ as well. However, if Player 2 does not win $\mathcal{G}(L(\mathcal{P}_d))$, then she might or might not win $\mathcal{G}(L(\mathcal{P}))$, i.e., the reduction is sound, but not complete, if \mathcal{P} is not good-for-games. The same holds true for Corollary 4.3.

While we only consider the realizability problem here, i.e., the problem of determining whether Player 2 wins the game, our proof of Theorem 4.2 can be extended to the synthesis problem, i.e., the problem of computing a winning strategy for Player 2, if she wins the game. Such a strategy can be finitely represented by a deterministic pushdown automaton with output (called pushdown transducers, or PDT) reading finite sequences over Σ_1 and outputting a single letter from Σ_2 . These are efficiently computable for Gale-Stewart games with ω -DCFL winning conditions [Fri10, Wal01]. Hence, one can compute a winning strategy for Player 2 in $\mathcal{G}(L(\mathcal{P}_d))$ and then apply the transformation described in the second part of the proof of Theorem 4.2, which is implementable by deterministic pushdown transducers.

Theorem 4.4. *Let \mathcal{P} be an ω -GFG-PDA. If Player 2 wins $\mathcal{G}(L(\mathcal{P}))$, then she has a winning strategy that is implemented by a PDT. Furthermore, such a PDT can be computed in exponential time in $|\mathcal{P}|$.*

The details of this straightforward, but slightly tedious, construction are presented in the appendix.

On the other hand, there are games with ω -GFG-CFL winning condition that are won by Player 1, but not with a winning strategy implementable by pushdown transducers. Hence,

the situation between the players is asymmetric. Again, the example proving this claim is presented in the appendix.

5. CLOSURE PROPERTIES

In Section 3, we have shown that ω -GFG-CFL is a new subclass of ω -contextfree languages. Here, we study the closure properties of this class, which differ considerably from those of ω -DCFL and ω -CFL.

We say that a class \mathcal{L} of ω -languages is closed under (ε -free) homomorphisms if $\{f(w) \mid w \in L\}$ is in \mathcal{L} for every $L \in \mathcal{L}$ and every $f: \Sigma \rightarrow (\Sigma')^+$. Here, we disallow ε in the image of f to ensure that $f(w)$ is infinite.

Theorem 5.1. *ω -GFG-CFL is not closed under union, intersection, complementation, set difference, nor homomorphism.*

Proof. Union: As ω -DCFL contains both $L_1 = \{(a\#)^n(b\#)^n\#^\omega \mid n \geq 1\}$ and $L_2 = \{(a\#)^n(b\#)^{2n}\#^\omega \mid n \geq 1\}$ whose union is not in ω -GFG-CFL (Theorem 3.7), Proposition 3.2 implies that ω -GFG-CFL is not closed under union.

Intersection: As ω -DCFL contains $L_1 = \{a^n b^n a^* b^\omega \mid n \geq 1\}$ and $L_2 = \{a^* b^n a^n b^\omega \mid n \geq 1\}$ whose intersection $L_1 \cap L_2 = \{a^n b^n a^n b^\omega \mid n \geq 1\}$ is not even in ω -CFL [CG77b, Proposition 1.3], Proposition 3.2 implies that ω -GFG-CFL is not closed under intersection.

Complementation: We show that the complement of the language $L_{ss} \in \omega$ -GFG-CFL used in Section 3 to separate ω -DCFL and ω -GFG-CFL is not even in ω -CFL. A word w is not in L_{ss} if $\pi_i(w)$ does not have a safe suffix for both $i \in \{1, 2\}$, which due to Remark 3.3 is equivalent to

$$\liminf_{n \rightarrow \infty} EL(\pi_i(w(0) \cdots w(n))) = -\infty$$

for both i .

Towards a contradiction, assume there is an ω -PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, \Omega)$ with $L(\mathcal{P}) = \Sigma^\omega \setminus L_{ss}$. As in the proof of Lemma 3.5, define $x_1 = \begin{pmatrix} + \\ 0 \end{pmatrix} \begin{pmatrix} + \\ - \end{pmatrix}$ and $x_2 = \begin{pmatrix} 0 \\ + \end{pmatrix} \begin{pmatrix} - \\ + \end{pmatrix}$. Recall that every infix of length at least 3 of a word built by concatenating copies of the x_i has a strictly positive energy level in one component.

Again, we define

$$w_{ss} = x_1 (x_2)^3 (x_1)^7 (x_2)^{15} (x_1)^{31} (x_2)^{63} \cdots$$

which satisfies

$$EL(\pi_1(x_1 (x_2)^3 \cdots (x_2)^{2^{2j}-1})) = -j$$

for every $j > 1$ and

$$EL(\pi_2(x_1 (x_2)^3 \cdots (x_1)^{2^{2j-1}-1})) = -j$$

for every $j > 0$ (see Page 9). Hence, $w_{ss} \in \Sigma^\omega \setminus L_{ss}$, i.e., there is an accepting run $\rho = c_0 \tau_0 c_1 \tau_1 c_2 \tau_2 \cdots$ of \mathcal{P} on w .

Recall that a step of ρ is a position n such that $\text{sh}(c_n) \leq \text{sh}(c_{n+j})$ for all $j \geq 0$. Every infinite run has infinitely many steps. Hence, we can find two steps $s < s'$ satisfying the following properties:

- (1) There is a state $q \in Q$ and a stack symbol $X \in \Gamma_\perp$ such that $c_s = (q, \gamma X)$ and $c_{s'} = (q, \gamma' X)$ for some γ, γ' , i.e., both configurations have the same state and topmost stack symbol.

- (2) The maximal color labeling the sequence $\tau_s \cdots \tau_{s'-1}$ of transitions leading from c_s to $c_{s'}$ is even.
- (3) The sequence $\tau_s \cdots \tau_{s'-1}$ processes an infix v of w with $EL(\pi_i(v)) > 0$, for some $i \in \{1, 2\}$ (recall that every sufficiently long infix of a word constructed from the x_i has positive energy level in one component).

Consider the sequence $\tau_0 \cdots \tau_{s-1}(\tau_s \cdots \tau_{s'-1})^\omega$ of transitions. Due to the first property, it induces a run ρ' of \mathcal{P} , which is accepting due to the second property. Finally, due to the third property, ρ' processes a word with suffix v^ω . Such a word has a safe suffix in component i , as $EL(\pi_i(v)) > 0$.

Hence, we have constructed a word w in L_{ss} such that there is an accepting run of \mathcal{P} on w , i.e., we have derived a contradiction to $L(\mathcal{P}) = \Sigma^\omega \setminus L_{ss}$. As we have picked \mathcal{P} arbitrarily, we have shown that $\Sigma^\omega \setminus L_{ss}$ is not in ω -CFL. Thus, due to Proposition 3.2 and Lemma 3.4, ω -GFG-CFL is not closed under complementation.

Set difference: As Σ^ω is in ω -DCFL $\subseteq \omega$ -GFG-CFL for every alphabet Σ , ω -GFG-CFL cannot be closed under set difference, as complementation is set difference with Σ^ω .

Homomorphism: As ω -DCFL contains

$$L = \left\{ \left[\begin{pmatrix} a \\ 1 \end{pmatrix} \begin{pmatrix} \# \\ 1 \end{pmatrix} \right]^n \left[\begin{pmatrix} b \\ 1 \end{pmatrix} \begin{pmatrix} \# \\ 1 \end{pmatrix} \right]^n \begin{pmatrix} \# \\ 1 \end{pmatrix}^\omega \mid n \geq 1 \right\} \cup \left\{ \left[\begin{pmatrix} a \\ 2 \end{pmatrix} \begin{pmatrix} \# \\ 2 \end{pmatrix} \right]^n \left[\begin{pmatrix} b \\ 2 \end{pmatrix} \begin{pmatrix} \# \\ 2 \end{pmatrix} \right]^{2n} \begin{pmatrix} \# \\ 2 \end{pmatrix}^\omega \mid n \geq 1 \right\}$$

whose projection (which is a homomorphism)

$$\pi_1(L) = \{(a\#)^n(b\#)^n\#^\omega \mid n \geq 1\} \cup \{(a\#)^n(b\#)^{2n}\#^\omega \mid n \geq 1\}$$

is not in ω -GFG-CFL (Theorem 3.7), Proposition 3.2 implies that ω -GFG-CFL is not closed under homomorphisms. \square

As we only used languages in ω -DCFL $\subseteq \omega$ -GFG-CFL to witness the failure of closure under intersection, union, and set difference, we obtain the following corollary.

Corollary 5.2. *ω -GFG-CFL is not closed under union, intersection, and set difference with languages in ω -DCFL.*

Finally, using standard arguments one can show that closure under these operations with ω -regular languages holds, as it does for ω -DCFL and ω -CFL.

Theorem 5.3. *If $L \in \omega$ -GFG-CFL and R is ω -regular, then $L \cap R$, $L \cup R$, and $L \setminus R$ are in ω -GFG-CFL as well.*

Proof. Let $L = L(\mathcal{P})$ for some ω -GFG-PDA \mathcal{P} and R be ω -regular, i.e., $R = L(\mathcal{A})$ for some deterministic parity automaton \mathcal{A} (see, e.g., [GTW02] for definitions). Furthermore, let $\mathcal{P} \times \mathcal{A}$ be the product automaton of these two automata, which is again an ω -PDA that simulates a run of \mathcal{P} and the unique run of \mathcal{A} simultaneously.

Using a detour via the Muller acceptance condition and the LAR construction (see for example [GTW02]), one can turn $\mathcal{P} \times \mathcal{A}$ into automata $(\mathcal{P} \times \mathcal{A})_\cap$, $(\mathcal{P} \times \mathcal{A})_\cup$, and $(\mathcal{P} \times \mathcal{A})_\setminus$ such that the following holds true:

- A run of $(\mathcal{P} \times \mathcal{A})_\cap$ is accepting if the simulated run of \mathcal{P} and the simulated run of \mathcal{A} are accepting.

- A run of $(\mathcal{P} \times \mathcal{A})_{\cup}$ is accepting if either the simulated run of \mathcal{P} or the simulated run of \mathcal{A} is accepting.
- A run of $(\mathcal{P} \times \mathcal{A})_{\setminus}$ is accepting if the simulated run of \mathcal{P} is accepting and the simulated run of \mathcal{A} is not accepting.

All three automata can be shown to be good-for-games, as only the nondeterminism of \mathcal{P} has to be resolved: A resolver for \mathcal{P} can easily be turned into one for the three automata that just ignores the additional inputs stemming from taking the product (note that this crucially depends on \mathcal{A} and the LAR memory being deterministic). \square

Note that $R \setminus L$ is not necessarily in ω -GFG-CFL (not even in ω -CFL) if R is ω -regular and L is in ω -GFG-CFL.

6. DECIDING GOOD-FOR-GAMENESS

In this section, we show that deciding good-for-gameness is, unfortunately, undecidable, both for automata and for languages. These results contrast with good-for-gameness being decidable for parity automata [KM15], even in polynomial time for Büchi and co-Büchi automata [KM15, BK18], and every ω -regular language being good-for-games, as deterministic parity automata recognize all ω -regular languages.

To show these undecidability results, we introduce some additional notation. If \mathcal{P} is an ω -PDA and q one of its states, then we write $L(\mathcal{P}, q)$ for the language accepted by the automaton obtained from \mathcal{P} by replacing its initial state with q .

Theorem 6.1. *The following problems are undecidable:*

- (1) *Given an ω -PDA \mathcal{P} , is \mathcal{P} good-for-games?*
- (2) *Given an ω -PDA \mathcal{P} , is $L(\mathcal{P}) \in \omega$ -GFG-CFL?*

Proof. Both proofs proceed by reduction from an undecidable problem for PDA over finite words (see, e.g., [HMU07]). Such an automaton has the same structure as an ω -PDA, but the coloring Ω is replaced by a set F of accepting states. A finite run is accepting, if it ends in an accepting state.

We consider the following two problems:

- DPDA inclusion: Given DPDA's \mathcal{D}_1 and \mathcal{D}_2 , is $L(\mathcal{D}_1) \subseteq L(\mathcal{D}_2)$?
- PDA universality: Given a PDA \mathcal{P} over Σ , is $L(\mathcal{P}) = \Sigma^*$?

1.) We reduce the inclusion problem for DPDA over finite words to deciding whether an ω -PDA is good-for-games. Since the inclusion problem is undecidable [Val73], so is deciding whether an ω -PDA is good-for-games.

Given DPDA \mathcal{D}_1 and \mathcal{D}_2 over Σ^* , we first define infinitary versions of \mathcal{D}_1 and \mathcal{D}_2 : Let \mathcal{P}_1 and \mathcal{P}_2 be ω -DPDA over $(\Sigma \cup \{\#\})^\omega$, where $\# \notin \Sigma$, that are identical to \mathcal{D}_1 and \mathcal{D}_2 respectively, except with additional $\#$ -transitions from states that are accepting in \mathcal{D}_1 and \mathcal{D}_2 to accepting sinks; all other states are made rejecting. Then $L(\mathcal{P}_i)$ consists of words of the form $v\#w$ where $v \in L(\mathcal{D}_i)$ and $w \in (\Sigma \cup \{\#\})^\omega$. We have $L(\mathcal{P}_1) \subseteq L(\mathcal{P}_2)$ if and only if $L(\mathcal{D}_1) \subseteq L(\mathcal{D}_2)$.

Consider \mathcal{P} , an ω -PDA over $\Sigma \cup \{\#, \$\}$ built as follows (see Figure 6): A fresh initial state q_I has $\$$ -transitions to fresh states q_1 and q_2 ; from q_1 there is an ε -transition to the initial state of \mathcal{P}_1 , and from q_2 there is a $\$$ -transition to an accepting sink q_s and an ε -transition to the initial state of \mathcal{P}_2 . None of these transitions manipulate the stack. We show that \mathcal{P} is good-for-games if and only if $L(\mathcal{D}_1) \subseteq L(\mathcal{D}_2)$.




Figure 6: The construction of \mathcal{P} from \mathcal{P}_1 and \mathcal{P}_2 . All new transitions do not manipulate the stack, i.e., if \mathcal{P}_i is entered via the ε -transition, then the initial configuration of \mathcal{P}_i is reached.

If $L(\mathcal{D}_1) \subseteq L(\mathcal{D}_2)$, then $L(\mathcal{P}_1) \subseteq L(\mathcal{P}_2)$ and $L(\mathcal{P}, q_1) \subseteq L(\mathcal{P}, q_2)$. Let r be defined such that

- $r(\varepsilon, \$) = (q_I, \perp, \$, q_2, \perp)$,
- $r((q_I, \perp, \$, q_2, \perp), \$) = (q_2, \perp, \$, q_s, \perp)$, and
- $r((q_I, \perp, \$, q_2, \perp), a) = (q_2, \perp, \$, q_I^2, \perp)$ where $a \neq \$$ and q_I^2 denotes the initial state of \mathcal{P}_2 , i.e., r produces a run prefix that reaches either the accepting sink q_s or the initial state q_I^2 of \mathcal{P}_2 , which is deterministic. Hence, in both cases, there is no further nondeterminism to resolve.

We claim that r is a resolver. Let $w \in L(\mathcal{P})$, i.e., either $w = \$w'$ with

$$w' \in L(\mathcal{P}, q_1) \cup L(\mathcal{P}, q_2) \subseteq L(\mathcal{P}, q_2),$$

i.e., the second letter of w is not equal to $\$$, or $w = \$\w' with $w' \in (\Sigma \cup \{\#, \$\})^\omega$. In both cases, the run of \mathcal{P} on w induced by r is accepting.

Conversely, if there is a word $v \in L(\mathcal{D}_1) \setminus L(\mathcal{D}_2)$, there is no resolver for \mathcal{P} . Indeed, if a resolver r chooses the transition $(q_I, \perp, \$, q_1, \perp)$ to q_1 to begin with, it is not a resolver since $\$^\omega \in L(\mathcal{P})$, but $\$^\omega \notin L(\mathcal{P}, q_1)$. Finally, if r chooses the transition $(q_I, \perp, \$, q_2, \perp)$ to q_2 to begin with, then r is not a resolver either since $\$v\#^\omega \in L(\mathcal{P})$, but $v\#^\omega \notin L(\mathcal{P}, q_2)$.

2.) We proceed by a reduction from the universality problem for PDA over finite words, which is undecidable [HMU07]. Namely, given a PDA \mathcal{F} over Σ we build an ω -PDA \mathcal{P} over $\Sigma_\# \times \{a, b, \#\}$ with $\Sigma_\# = \Sigma \cup \{\#\}$ such that $L(\mathcal{P}) \in \omega\text{-GFG-CFL}$ if and only if $L(\mathcal{F})$ is universal.

First, note that

$$L_1 = \{v\#w \mid v \in L(\mathcal{F}) \text{ and } w \in \Sigma_\#^\omega\} \cup \Sigma^\omega$$

is universal if and only if $L(\mathcal{F})$ is universal. Furthermore, \mathcal{F} can easily be turned into an ω -PDA recognizing L_1 . So, we can construct from \mathcal{F} an ω -PDA \mathcal{P} recognizing the language

$$L(\mathcal{P}) = \{w \in (\Sigma_\# \times \{a, b, \#\})^\omega \mid \pi_1(w) \in L_1 \text{ or } \pi_2(w) \in L_2\},$$

where L_2 is the ω -language from the proof of Theorem 3.7, which is not in $\omega\text{-GFG-CFL}$ but in $\omega\text{-CFL}$. We claim that \mathcal{P} has the desired property. Trivially, if $L(\mathcal{F})$ is universal, then so is $L(\mathcal{P})$, which implies that $L(\mathcal{P})$ is in $\omega\text{-GFG-CFL}$.

Now assume that $L(\mathcal{F})$ is not universal which is witnessed by some word $v\#^\omega \notin L_1$, i.e., such that $v \notin L(\mathcal{F})$. Towards a contradiction, assume that there is an $\omega\text{-GFG-PDA}$ \mathcal{R} with $L(\mathcal{R}) = L(\mathcal{P})$, say with resolver r . We turn \mathcal{R} into another $\omega\text{-GFG-PDA}$ \mathcal{R}' recognizing the language

$$\left\{ w \in \{a, b, \#\}^\omega \mid \begin{pmatrix} v\#^\omega \\ w \end{pmatrix} \in L(\mathcal{R}) \right\},$$

which yields the desired contradiction, as this language is equal to L_2 , which is not in ω -GFG-CFL.

To this end, we equip \mathcal{R} with a counter ranging over the positions of $v\#$ to simulate a run of \mathcal{R} on input $(\frac{v\#^\omega}{w})$ when given the input w . As the counter behaves deterministically, no new nondeterminism is introduced when constructing \mathcal{R}' from \mathcal{R} . Hence, r can be turned into a resolver for \mathcal{R}' , which is therefore good-for-games, a contradiction. \square

7. RESOURCE-BOUNDED RESOLVERS

As defined Section 3, a resolver is an arbitrary function, potentially even an uncomputable one, mapping a run prefix ρ and the next letter to be processed to a transition that is enabled in the last configuration of ρ . Here, we study resource-bounded resolvers.

For ω -regular GFG automata, finite-state resolvers are always sufficient [HP06]. Unsurprisingly, finite-state resolvers for ω -GFG-PDA are too weak. In fact, we show that ω -GFG-PDA with finite-state resolvers are determinizable.

On the other hand, it is tempting to conjecture that every ω -GFG-PDA has a resolver that can be implemented by a pushdown automaton with output. However, we disprove this claim, i.e., even stronger devices are necessary in general.

Before we present our results, let us discuss one technical detail. Consider an ω -GFG-PDA \mathcal{P} and a resolver r for \mathcal{P} . As r always has to return an enabled transition, which depends on the current state and the current top stack symbol of \mathcal{P} , a machine computing r would need to keep track of this information. However, this requires, in general, keeping track of the full stack of \mathcal{P} . This is beyond the capabilities of finite-state machines and blocks the stack of a pushdown machine implementing r from performing other computations. In order to better capture the complexity of resolving nondeterminism without involving the complexity of tracking enabled transitions, we equip the machine computing r with an oracle providing the current top stack symbol of \mathcal{P} .

We begin by considering finite-state resolvers. To this end, fix an ω -GFG-PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, \Omega)$. A Moore machine $\mathcal{M} = (M, m_I, \delta, \lambda)$ for \mathcal{P} consists of a finite set M of states with an initial state $m_I \in M$, a transition function $\delta: M \times \Delta \rightarrow M$, and an output function $\lambda: M \times \Sigma \times \Gamma_\perp \rightarrow \Delta$. As usual, we extend the transition function to $\delta: \Delta^* \rightarrow M$ by defining $\delta(\varepsilon) = m_I$ and $\delta(\tau_0 \cdots \tau_{n-1} \tau_n) = \delta(\delta(\tau_0 \cdots \tau_{n-1}), \tau_n)$.

Let $\tau_0 \cdots \tau_n \in \Delta^*$ and $a \in \Sigma$. If the sequence $\tau_0 \cdots \tau_n$ induces a run prefix of \mathcal{P} , then let X be the top stack symbol of its last configuration. In this case, we define $f_{\mathcal{M}}(\tau_0 \cdots \tau_n, a) = \lambda(\delta(\tau_0 \cdots \tau_n), a, X)$. If $\tau_0 \cdots \tau_n$ does not induce a run prefix, then we define $f_{\mathcal{M}}(\tau_0 \cdots \tau_n, a)$ arbitrarily. We say that a resolver r for \mathcal{P} is a finite-state resolver, if there is some Moore machine \mathcal{M} for \mathcal{P} such that $r = f_{\mathcal{M}}$. Note that the Moore machine \mathcal{M} on input $\tau_0 \cdots \tau_n$ has to output a transition that is enabled in the last configuration of the run prefix induced by $\tau_0 \cdots \tau_n$.

Example 7.1. Recall the ω -GFG-PDA \mathcal{P} in Example 2.2 on Page 6 recognizing the language

$$\{ac^n d^n \#^\omega \mid n \geq 1\} \cup \{bc^n d^{2n} \#^\omega \mid n \geq 1\}.$$

It has a finite-state resolver, which is depicted in Figure 7. Intuitively, it remembers the first letter of the input word processed, which suffices to resolve the nondeterminism at state q_1 of \mathcal{P} . This is the only nondeterministic choice to make, the rest of the run can be determined by keeping track of the current state of \mathcal{P} and using the information on the stack.




Figure 7: A finite-state resolver for the ω -GFG-PDA \mathcal{P} from Example 2.2. Here, X is again an arbitrary stack symbol of \mathcal{P} (not including \perp). All missing transitions lead to a rejecting sink (not depicted) and only relevant outputs are given (in light gray boxes near states), all others can be picked arbitrarily.

Recall that the language recognized by the ω -GFG-PDA in Example 2.2 is in ω -DCFL. This is no coincidence.

Theorem 7.2. *If an ω -GFG-PDA \mathcal{P} has a finite-state resolver, then $L(\mathcal{P}) \in \omega$ -DCFL.*

Proof. Given an ω -GFG-PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, \Omega)$ and a Moore machine $\mathcal{M} = (M, m_I, \delta, \lambda)$ for \mathcal{P} implementing a resolver, we construct an ω -DPDA \mathcal{D} with $L(\mathcal{D}) = L(\mathcal{P})$. The set of states of \mathcal{D} is $(Q \times M) \cup (Q \times M \times \Sigma)$ with initial state (q_I, m_I) , and its stack alphabet is Γ . Intuitively, \mathcal{D} simulates the run constructed by $r_{\mathcal{M}}$. To this end, it has to keep track of a state of \mathcal{P} and a state of \mathcal{M} . In state (q, m) , \mathcal{D} deterministically reads the next input letter to be processed and stores it by moving to state (q, m, a) . Now, this information, together with the top stack symbol X of the current configuration, suffices to determine the transition $\lambda(m, a, X)$, which is then simulated by updating the state of \mathcal{P} , the state of \mathcal{M} , and the stack.

Formally define the transitions and their colors as follows:

- For every state $(q, m) \in Q \times M$, every $X \in \Gamma_{\perp}$, and every $a \in \Sigma$, \mathcal{D} has the transition $((q, m), X, a, (q, m, a), \gamma)$ of color $\min \Omega(\Delta)$.
- For every state $(q, m, a) \in Q \times M \times \Sigma$ and every $X \in \Gamma_{\perp}$ such that $\lambda(m, a, X)$ is an ε -transition of the form $\tau = (q', \varepsilon, q', \gamma)$ for some q' and some γ , \mathcal{D} has the transition $((q, m, a), X, \varepsilon, (q', \delta(m, \tau), a), \gamma)$ of color $\Omega(\tau)$.

- For every state $(q, m, a) \in Q \times M \times \Sigma$ and every $X \in \Gamma_\perp$ such that $\lambda(m, a, X)$ is an a -transition of the form $\tau = (q, X, a, q', \gamma)$ for some q' and some γ , \mathcal{D} has the transition $((q, m, a), X, \varepsilon, (q', \delta(m, \tau)), \gamma)$ of color $\Omega(\tau)$.

By construction, there is exactly one enabled a -transition in every configuration with state in $Q \times M$ and there is at most one enabled ε -transition in each configuration with state in $Q \times M \times \Sigma$. Hence, \mathcal{D} is indeed deterministic. Finally, an induction shows that $L(\mathcal{D})$ is equal to $L(\mathcal{P})$ and therefore $L(\mathcal{P}) \in \omega\text{-DCFL}$. \square

Next, we consider resolvers implemented by pushdown automata with output. Formally, a pushdown transducer (PDT) for an ω -GFG-PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, \Omega)$ is a tuple $\mathcal{T} = (\bar{Q}, \Delta, \bar{\Gamma}, \bar{q}_I, \bar{\Delta}, \bar{\lambda})$ such that $(\bar{Q}, \Delta, \bar{\Gamma}, \bar{q}_I, \bar{\Delta})$ is a deterministic PDA (without coloring) processing transitions of \mathcal{P} and $\bar{\lambda}: \bar{Q} \times \Sigma \times \Gamma_\perp \rightarrow \Delta$ is a (partial) output function.

A run of \mathcal{T} on an input $v \in \Delta^*$ is a sequence $c_0 \tau_0 c_1 \tau_1 \cdots \tau_n c_n$ such that $c_0 = (\bar{q}_I, \perp)$ and $c_i \xrightarrow{\tau_i} c_{i+1}$ for every $i < n$, $\ell(\tau_0 \cdots \tau_n) = v$, and c_n has no enabled ε -transitions. Due to determinism and the requirement that the last configuration has no enabled ε -transition, \mathcal{T} has at most one run on every input v .

Let $\tau_0 \cdots \tau_n \in \Delta^*$ and $a \in \Sigma$. If the sequence $\tau_0 \cdots \tau_n$ induces a run prefix of \mathcal{P} , then let X be the top stack symbol of its last configuration. In this case, we define $f_{\mathcal{T}}(\tau_0 \cdots \tau_n, a) = \bar{\lambda}(\bar{q}, a, X)$, where \bar{q} is the state of the last configuration of the run of \mathcal{T} on $\tau_0 \cdots \tau_n$. If $\tau_0 \cdots \tau_n$ does not induce a run prefix of \mathcal{P} , then we define $f_{\mathcal{T}}(\tau_0 \cdots \tau_n, a)$ arbitrarily. Again, we say that a resolver r for \mathcal{P} is a pushdown resolver, if there is some PDT \mathcal{T} for \mathcal{P} such that $r = f_{\mathcal{T}}$.

Theorem 7.3. *There is no pushdown resolver for the ω -GFG-PDA \mathcal{P} in Figure 4.*

Proof. We will proceed by contradiction. Assume r is a pushdown resolver, implemented by some PDT \mathcal{T} , for the ω -GFG-PDA \mathcal{P} in Figure 4 on Page 8, which recognises L_{ss} . Due to Remark 3.1, every finite word has a run prefix processing it that is consistent with r . Furthermore, since \mathcal{P} does not have any ε -transitions this run is unique.

Define the word w by setting $w(0)$ to $\binom{-}{+}$ and $w(n)$ for $n > 0$ to $\binom{-}{+}$ if the run induced by r over $w(0) \cdots w(n-1)$ ends in state 1, and to $\binom{+}{-}$ otherwise. Observe that the run induced by r over this word always has an empty stack and therefore the r -consistent run ρ over w is rejecting as it only uses transitions with color 1. Since r is a resolver, it follows that $w \notin L_{ss}$. Furthermore, since the stack remains empty throughout ρ , every position is a step in ρ , and the top stack symbol is always \perp (see Page 9 for the definition of steps).

Now, consider the run ρ' of \mathcal{T} generating the run ρ processing w in \mathcal{P} . Since \mathcal{T} is a PDT, there are infinitely many steps along ρ' . Observe that if two positions s and s' are steps in ρ' and agree on the top stack symbol in ρ' and the state in both ρ and ρ' , then the words $w(s)w(s+1)w(s+2) \cdots$ and $w(s')w(s'+1)w(s'+2) \cdots$ are identical: this is because $w(n+1)$ only depends on the transition chosen by \mathcal{T} after generating the run over $w(0) \cdots w(n)$, but \mathcal{T} is deterministic and at steps its behaviour does not depend on the stack configurations below the top stack symbols. Since both the stack alphabet of \mathcal{T} and state spaces of both \mathcal{P} and \mathcal{T} are finite, by the pigeonhole principle, such positions s and s' must exist. It follows that w is ultimately periodic: $w = uv^\omega$ for some nonempty v .

We now argue that $w \in L_{ss}$. If v has more occurrences of $\binom{+}{-}$ than $\binom{-}{+}$, then the first component has a safe suffix; similarly, if v has more occurrences of $\binom{-}{+}$, then the second component has a safe suffix; finally, if there are as many of both letters, then both

components have a safe suffix. In all cases, $w \in L_{ss}$, a contradiction. We conclude that there is no pushdown resolver for \mathcal{P} . \square

This is perhaps surprising given Theorem 4.4 (in the appendix), which states that Player 2 has winning strategies implementable by pushdown transducers in Gale-Stewart games with ω -GFG-CFL winning conditions. In particular, a consequence of Theorem 4.4 is that a universal ω -GFG-PDA \mathcal{P} has a pushdown resolver. This can be seen by applying the construction presented in Section 4 and in Appendix B to construct a winning strategy in the Gale-Stewart game with winning condition $\{(\frac{w}{\# \omega}) \mid w \in L(\mathcal{P})\}$, which is essentially a pushdown resolver for \mathcal{P} .

8. THE PARITY INDEX HIERARCHY

In this section we study how ω -GFG-PDA with different acceptance conditions compare in their expressivity. In particular, we establish that the classes of ω -GFG-PDA with a fixed number of colors form an infinite hierarchy: for each $n > 1$, there is a language recognized by an ω -GFG-PDA with n colors, but not recognized by an ω -GFG-PDA with $n - 1$ colors. This is also the case for ω -DPDA [Löd15] while ω -PDA with Büchi acceptance (i.e., with colors 1 and 2) suffice to recognize all ω -CFL [CG77a].

Lemma 8.1. *For each $n > 1$, there is a language L_n recognized by an ω -GFG-PDA with n colors but not recognized by any ω -GFG-PDA with $n - 1$ colors.*

Proof. For each $n > 1$, let $D_n = \{1, \dots, n\}$ and let $L_n \subseteq D_n^\omega$ be the language of ω -words satisfying the parity condition, that is, in which the highest color that occurs infinitely often is even. L_n is recognized by a deterministic parity automaton with n colors that upon reading a letter p visits a state of color p . We can view this parity automaton as an ω -GFG-PDA (even an ω -DPDA) that does not use its stack.

We now show that L_n is not recognized by an ω -GFG-PDA with only $n - 1$ colors. Towards a contradiction, assume there is such an ω -GFG-PDA \mathcal{P} with resolver r . The intuition of the argument is that we can build an infinite word w of which the sequence of colors is (roughly) the sequence of colors on its r -consistent run, incremented by one, leading to a contradiction. Since we allow ε -transitions, we use the maximal color over a sequence of transitions to choose the next letter of w , rather than just the last transition. Without loss of generality, we can assume that the set of colors of \mathcal{P} is either $\{1, \dots, n - 1\}$ or $\{0, \dots, n - 2\}$. Consider the word $w(0)w(1)w(2)\dots$ over D_n built as follows:

- $w(0) = 1$.
- For $i \geq 1$, first define ρ_j for $j \geq i$ to be the shortest r -consistent run processing $w(0)\dots w(j - 1)$. Then, $w(i) = p + 1$ where p is the highest color in the suffix of ρ_i starting after ρ_{i-1} (or at the beginning of the run if $i = 1$). In other words, p is the maximal color in the run ρ_i on $w(0)\dots w(i - 1)$ after the transition processing $w(i - 2)$.

We now argue that the word w defined as above induces an accepting r -consistent run if and only if it is not in L_n . To this end, define p to be the highest color occurring infinitely often along the r -consistent run processing w , which exists due to Remark 3.1. By construction, the highest color occurring infinitely often on w is $p + 1$. Therefore, the run is accepting if and only if w is not in L_n , contradicting that \mathcal{P} recognizes L_n . \square

The languages L_n used in the above proof are ω -regular, and recognized by deterministic parity automata with n colors, i.e., with the same number of colors as an ω -GFG-PDA for

L_n . However, by combining the languages L_n with L_{ss} , it is not hard to find a family of languages that requires both n colors and an ω -GFG-PDA. Indeed, the language

$$\left\{ \begin{pmatrix} w_1 \\ w_2 \end{pmatrix} \mid w_1 \in L_{ss} \text{ and } w_2 \in L_n \right\}$$

is recognized by an ω -GFG-PDA with n colors, obtained by taking the product of an ω -GFG-PDA for L_{ss} with a deterministic parity automaton for L_n . However, it is not recognized by an ω -DPDA, due to L_{ss} , nor by an ω -GFG-PDA with fewer than n colors, due to L_n .

9. COMPARISON TO VISIBLY PUSHDOWN LANGUAGES

In this section, we compare ω -GFG-CFL to another important subclass of ω -CFL, the class of visibly pushdown languages [AM04], for which solving games is decidable as well [LMS04].

Visibly pushdown automata are defined with respect to a partition $\tilde{\Sigma} = (\Sigma_c, \Sigma_r, \Sigma_s)$ of the input alphabet and have to satisfy the following conditions:

- A letter $a \in \Sigma_c$ is only processed by transitions of the form (q, X, a, q', XY) with $X \in \Gamma_\perp$, i.e., some stack symbol Y is pushed onto the stack.
- A letter $a \in \Sigma_r$ is only processed by transitions of the form $(q, X, a, q', \varepsilon)$ with $X \neq \perp$ or (q, \perp, a, q', \perp) , i.e., the topmost stack symbol is removed, or if the stack is empty, it is left unchanged.
- A letter $a \in \Sigma_s$ is only processed by transitions of the form (q, X, a, q', X) with $X \in \Gamma_\perp$, i.e., the stack is left unchanged.
- There are no ε -transitions.

Intuitively, the stack height of the last configuration of a run processing some $v \in (\Sigma_c \cup \Sigma_r \cup \Sigma_s)^*$ only depends on v .

A language $L \subseteq \Sigma^\omega$ is in ω -VPL if there is a partition $\tilde{\Sigma}$ of Σ such that there is a nondeterministic ω -visibly pushdown automaton \mathcal{P} recognizing L with respect to $\tilde{\Sigma}$.

Theorem 9.1. *ω -GFG-CFL and ω -VPL are incomparable with respect to inclusion.*

Proof. The language $L_{ss} \in \omega$ -GFG-CFL used in Section 3 to separate ω -GFG-CFL and ω -DCFL is not in ω -VPL, as ω -VPL $\subseteq \omega$ -CFL is closed under complementation [AM04] while the complement of L_{ss} is not in ω -CFL (Theorem 5.1). Thus, ω -GFG-CFL is not included in ω -VPL.

For the other non-inclusion, let $\Sigma = \{+, -\}$ and define the value⁸ $\text{val}(v) \in \mathbb{N}$ of a finite word $v \in \Sigma^*$ inductively as $\text{val}(\varepsilon) = 0$ as well as $\text{val}(v+) = \text{val}(v) + 1$ and $\text{val}(v-) = \max\{0, \text{val}(v) - 1\}$. This corresponds to the height of a stack after the sequence of push and pop operations described by v when $+$ is interpreted as a push and $-$ as a pop. We consider the language of such sequences that visit some value infinitely often. We show that

$$L_{\text{repbdd}} = \{w \in \Sigma^\omega \mid \text{there is an } s \in \mathbb{N} \text{ such that} \\ \text{val}(w(0) \cdots w(n)) = s \text{ for infinitely many } n\},$$

which is in ω -VPL [AM04], is not in ω -GFG-CFL.

To this end, fix an ω -GFG-PDA \mathcal{P} with resolver $r: \Delta^* \times \Sigma \rightarrow \Delta$. We will show that it does not recognize L_{repbdd} . We assume without loss of generality that all ε -transitions

⁸Alur and Madhusudan used the term *stack height* instead of *value*, but this is misleading here, since our automata are not necessarily visibly, i.e., the stack height of a run prefix on v might differ from $\text{val}(v)$.

have color 0 while all Σ -transitions have a nonzero color. This can be achieved by adding a component to \mathcal{P} 's states that accumulates the maximal color seen along a sequence of ε -transitions until a Σ -transition is used. As a consequence, a run of \mathcal{P} on some infinite input satisfies the acceptance condition if and only if the sequence of Σ -transitions satisfies the acceptance condition, i.e., the colors of ε -transitions are irrelevant and will be ignored in the following.

Given a run prefix $\rho \in \Delta^*$ and a letter $a \in \Sigma$, let $\text{ext}(\rho, a)$ be the unique extension of ρ induced by r when processing a . Formally, we define

$$\text{ext}(\rho, a) = \begin{cases} \rho \cdot r(\rho, a) & \text{if } \ell(r(\rho, a)) = a, \\ \text{ext}(\rho \cdot r(\rho, a), a) & \text{if } \ell(r(\rho, a)) = \varepsilon. \end{cases}$$

Due to Remark 3.1, if ρ is a run prefix that is consistent with r then $\text{ext}(\rho, a)$ is a finite extension of ρ .

We now build an ω -word w , letter by letter, based on how r resolves nondeterminism on the prefix built so far. The intuition is that whenever the run built by r sees an even color after a prefix v , the value of prefixes extending v remains above $\text{val}(v)$ until a larger odd color is seen, and then returns to $\text{val}(v)$ (unless a higher even color is seen in the meantime). Roughly, after an even color occurs in the run, the value of the word increases until a higher odd color occurs. Thus, if the maximal color that occurs infinitely often is even, i.e., the run is accepting, then the word will not be in the language. On the other hand, when an odd color p occurs, the value of the run decreases to just above the value of the word at the last even color higher than p . Hence if the maximal color that occurs infinitely often is odd, i.e., if the run is rejecting, then the value reaches the same level infinitely often, and the word is in the language. The result is that either w is in L_{repbdd} but rejected by \mathcal{P} , or w is not in L_{repbdd} but accepted by \mathcal{P} .

Formally, we inductively define an infinite sequence of sequences $\rho_n \in \Delta^*$, all ending in a Σ -transition. To start, we define $\rho_0 = \text{ext}(\varepsilon, +, +)$. To define ρ_{n+1} we have to consider several cases.

First, assume ρ_n ends with a $+$ -transition. Let p be the last and p' be the second-to-last nonzero color appearing in ρ_n (this is well-defined as ρ_0 contains two colors). If p is odd and $p \geq p'$, then we define $\rho_{n+1} = \text{ext}(\rho_n, -)$ (Case 1), otherwise, $\rho_{n+1} = \text{ext}(\rho_n, +)$ (Case 2).

Now, assume ρ_n ends with a $-$ -transition. Let $\text{sff}(\rho_n)$ be the suffix of ρ_n starting with the last $+$ -transition (this is well-defined, as ρ_0 contains $+$ -transitions). Furthermore, let $\text{prf}(\rho_n)$ be the prefix of ρ_n ending with the last transition having an even color that is at least as large as the maximal color labeling a transition in $\text{sff}(\rho_n)$. See Figure 8 for an illustration: $\text{prf}(\rho_n)$ is the prefix ending in the last transition that has an even color p that is at least as large as the colors p_0, \dots, p_j , the colors occurring in $\text{sff}(\rho_n)$. If there is no transition with such a color, then define $\text{prf}(\rho_n) = \varepsilon$. Note that $\text{sff}(\rho_n)$ and $\text{prf}(\rho_n)$ might overlap and that $\text{prf}(\rho_n) = \rho_n$ is possible if the last transition of ρ_n has an even color that is maximal among those in $\text{sff}(\rho_n)$. If

$$\text{val}(\ell(\rho_n)) > \text{val}(\ell(\text{prf}(\rho_n))) + 1,$$

then we define $\rho_{n+1} = \text{ext}(\rho_n, -)$ (Case 3), otherwise, $\rho_{n+1} = \text{ext}(\rho_n, +)$ (Case 4). Note that the even-numbered cases extend by a $+$, the odd-numbered cases by a $-$.

Now, let $\rho \in \Delta^\omega$ be the unique infinite sequence such that each ρ_n is a prefix of ρ , and let $w = \ell(\rho)$, which is an infinite word. Then, ρ induces a run of \mathcal{P} on w that is consistent with r .




Figure 8: Illustration of the definition of ρ_{n+1} in the case where ρ_n ends with a $--$ -transition. Each arrow depicts a transition of ρ_n , its color is depicted below the arrow, the letter it processes above (some transitions in $\text{sff}(\rho_n)$ may be ε -transitions, but only the first one is a $+$ -transition).

First, assume w is of the form $v^{+\omega}$. Then, from some point onwards, we only use Case 2 to extend ρ_n to ρ_{n+1} , i.e., if the last color of ρ_n is odd, then the second-to-last color is strictly larger. This implies that the maximal color occurring infinitely often in ρ is even. Thus, the run ρ is accepting although w is not in L_{repbdd} .

Now, assume w is of the form $v^{-\omega}$, i.e., $\text{val}(\ell(\rho_n)) = 0$ for almost all n . Then, from some point onwards, we only use Case 3 to extend ρ_n to ρ_{n+1} , i.e., we have

$$\text{val}(\ell(\rho_n)) > \text{val}(\ell(\text{prf}(\rho_n))) + 1$$

for almost all n . Combining both equations yields a contradiction, as $\text{val}(\ell(\text{prf}(\rho_n))) + 1$ is positive. Thus, w cannot have the form $v^{-\omega}$.

As a last case, assume w contains infinitely many $+$ and infinitely many $-$. First, we study the case where the maximal color occurring infinitely often in ρ , call it p , is odd. Let n_0 be such that the suffix of ρ obtained from removing ρ_{n_0} only contains colors that occur infinitely often and ρ_{n_0} contains at least one p that is not followed by a larger even color. Furthermore, let ρ^* be the longest prefix of ρ ending in a transition with an even color that is larger than p (which by construction is a prefix of ρ_{n_0}). We define $\rho^* = \varepsilon$ if there is no such color. Let $b = \text{val}(\rho^*)$. We claim that there are infinitely many n such that $\text{val}(\ell(\rho_n)) \leq b + 1$. Then, ρ is rejecting while w is in L_{repbdd} due to the pigeonhole principle.

To this end, let $n' > n_0$ such that $\rho_{n'}$ ends with a transition τ of color p . As there are infinitely many such n' , it suffices to show that there is a $n \geq n'$ such that $\text{val}(\ell(\rho_n)) \leq b + 1$. First, assume $\ell(\tau) = +$. By construction, the last nonzero color occurring before the final p is not larger than p . Hence, we have $\rho_{n'+1} = \text{ext}(\rho_{n'}, -)$ due to Case 1. Now, we either already have $\text{val}(\rho_{n'+1}) \leq b + 1$ or we apply Case 3 repeatedly until we have produced some ρ_n with $\text{val}(\rho_n) = b + 1$. The reason why Case 3 is always applicable is that the suffix $\text{sff}(\rho_{n'})$ for $n' + 1 \leq n'' \leq n$ always contains the last transition of $\rho_{n'}$, with color p . This in turn implies that the prefix $\text{prf}(\rho_{n''})$ is equal to ρ^* .

The case for $\ell(\tau) = -$ is similar: either we already have $\text{val}(\rho_{n'}) \leq b + 1$ or we apply Case 3 repeatedly until we have produced some ρ_n with $\text{val}(\rho_n) = b + 1$.

Finally, we consider the case where p , the maximal color occurring infinitely often in ρ , is even. We show that there are infinitely many n such that $\text{val}(\ell(\rho_{n'})) > \text{val}(\ell(\rho_n))$ for every $n' > n$. This implies that for every s there are only finitely many n such that $\text{val}(\ell(\rho_n)) = s$. As the prefixes of w are of the form $\ell(\rho_n)$ for $n \in \mathbb{N}$, we obtain that w is not in L_{repbdd} , even though the run of \mathcal{P} on w induced by ρ is accepting.

Let n_0 be such that the suffix of ρ obtained from removing ρ_{n_0} only contains colors that occur infinitely often and ρ_{n_0} contains a suffix ρ_s starting with a transition of color p such that ρ_s does not contain a larger color than p , but does contain a $+$ -transition. By definition,

the resulting suffix of ρ has infinitely many transitions of color p , all of which mark the end of some ρ_n with $n \geq n_0$. We show that each such n has the desired property.

By the choice of n_0 , the last transition of ρ_{n+1} is a $+$ -transition, no matter whether the last transition of ρ_n is a $+$ -transition or a $-$ -transition. If it is a $+$ -transition, then ρ_{n+1} is obtained by applying Case 2 to ρ_n , as the color of the last transition of ρ_n is even. On the other hand, if the last transition of ρ_n is a $-$ -transition, then ρ_{n+1} is obtained by applying Case 4 to ρ_n , as the prefix $\text{prf}(\rho_n)$ is equal to ρ_n by the fact that ρ_{n_0} contains an occurrence of a p that is not followed by a larger color, but by a $+$ -transition.

Furthermore, ρ_{n+2} is obtained by applying Case 2 to ρ_{n+1} , as the color of the last transition of ρ_{n+1} might be odd, but then it is strictly smaller than p , the second-to-last color in ρ_{n+1} . Therefore, the last transition of ρ_{n+1} is also a $+$ -transition.

Now, assume towards a contradiction that there is some $n' > n$ such that $\text{val}(\ell(\rho_{n'})) = \text{val}(\ell(\rho_n))$. Pick n' minimal with this property. Then, $n' > n + 2$ since the last transitions of both ρ_{n+1} and ρ_{n+2} are $+$ -transitions. Therefore, due to minimality, the last transition of $\rho_{n'}$ and the last transition of $\rho_{n'-1}$ are both $-$ -transitions. Hence, $\rho_{n'} = \text{ext}(\rho_{n'-1}, -)$ due to Case 3.

Now, consider the prefix $\text{prf}(\rho_{n'-1})$ in the application of Case 3 to $\rho_{n'-1}$. It is either equal to ρ_n (as its color is even and at least as large as all colors that may appear in the suffix $\text{sff}(\rho_{n'-1})$ of $\rho_{n'-1}$), or it is some extension of ρ_n . In both cases, we have $\text{val}(\ell(\rho_n)) \leq \text{val}(\ell(\text{prf}(\rho_{n'-1})))$ (in the latter due to the minimality of n'). Hence,

$$\begin{aligned} \text{val}(\ell(\rho_n)) &\leq \text{val}(\ell(\text{prf}(\rho_{n'} - 1))) < \text{val}(\ell(\rho_{n'-1})) - 1 \\ &= \text{val}(\ell(\rho_{n'})) = \text{val}(\ell(\rho_n)), \end{aligned}$$

which yields the desired contradiction. Here, the strict inequality follows from the definition of Case 3.

To conclude, in either of the cases we considered, the run ρ is accepting, but $w \notin L_{\text{repbdd}}$, or the run ρ induced by the resolver r is rejecting, but $w \in L_{\text{repbdd}}$. Hence, either \mathcal{P} does not recognize L_{repbdd} or r is not a resolver for \mathcal{P} , i.e., L_{repbdd} is not in ω -GFG-CFL. \square

10. CONCLUSION

We have introduced good-for-games ω -pushdown automata and proved that they recognize a novel class of ω -contextfree languages for which solving games (and synthesizing winning strategies) is decidable. Furthermore, we have studied (the mostly nonexistent) closure properties of the new class, proven that it is incomparable to ω -visibly pushdown languages, and that deciding good-for-gameness is undecidable for ω -pushdown automata and ω -contextfree languages. Finally, we proved the parity index hierarchy to be strict for ω -GFG-PDA and studied resource-bounded resolvers: finite-state resolvers only exist for ω -DCFL and even pushdown resolvers are not sufficient for ω -GFG-PDA.

We hope this paper is the catalyst for an in-depth study of good-for-games automata in settings where nondeterministic automata are more expressive than deterministic ones. But even for the setting of ω -contextfree languages considered here, we open many interesting directions for further research. Let us conclude by listing a few:

- Is the universality problem for ω -GFG-PDA EXPTIME-complete? Note that this problem is a promise problem with an undecidable promise, e.g., we only consider ω -GFG-PDA as input.

- Equivalence is decidable for ω -DPDA with weak acceptance conditions [LR12] (that is, each strongly connected component is either rejecting or accepting), but it is undecidable for ω -DPDA with Büchi or co-Büchi acceptance conditions [BGHH17]. Whether weak acceptance conditions make the problem decidable for ω -GFG-PDA is, to the best of our knowledge, open.
- Can ω -GFG-CFL be characterized by some extension of Monadic Second-order Logic? Such characterizations have been exhibited for contextfree languages of finite [LST94] and infinite words [DDK20] as well as for the class of visibly pushdown languages [AM04]. However, there is, to the best of our knowledge, no characterization of the deterministic contextfree languages, neither for finite nor infinite words.
- We have shown that there are ω -GFG-PDA that do not have pushdown resolvers, but left open whether for every ω -GFG-CFL there is an ω -GFG-PDA recognizing that language with a pushdown resolver. In fact it is also open whether every ω -GFG-PDA or every ω -GFG-CFL has a computable resolver.
- While here we focus on infinite words, good-for-games pushdown automata turn out to be more powerful than their deterministic counterparts, already over finite words. In [GJLZ21], we study their expressivity and succinctness over finite words, but some gaps remain open.
- Another interesting direction, proposed by one of the reviewers, is to further restrict the model of ω -GFG-PDA, by considering weaker classes of nondeterministic ω -PDA, e.g., unambiguous ones or one-counter automata. For finite words, unambiguity and good-for-gameness are incomparable [GJLZ21].

Acknowledgements. We thank Sven Schewe and Patrick Totzke for fruitful discussions leading to the results presented in Section 6.



This work is part of a project that receives funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 892704.

REFERENCES

- [AM04] Rajeev Alur and P. Madhusudan. Visibly pushdown languages. In László Babai, editor, *STOC 2004*, pages 202–211. ACM, 2004. doi:10.1145/1007352.1007390.
- [BGHH17] Stanislav Böhm, Stefan Göller, Simon Halfon, and Piotr Hofman. On Büchi One-Counter Automata. In Heribert Vollmer and Brigitte Vallée, editors, *STACS 2017*, volume 66 of *LIPIcs*, pages 14:1–14:13, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2017/7019>, doi:10.4230/LIPIcs.STACS.2017.14.
- [BK18] Marc Bagnol and Denis Kuperberg. Büchi Good-for-Games Automata Are Efficiently Recognizable. In Sumit Ganguly and Paritosh Pandya, editors, *FSTTCS 2018*, volume 122 of *LIPIcs*, pages 16:1–16:14. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018. URL: <http://drops.dagstuhl.de/opus/volltexte/2018/9915>, doi:10.4230/LIPIcs.FSTTCS.2018.16.
- [BKKM13] Udi Boker, Denis Kuperberg, Orna Kupferman, and Michał Skrzypczak. Nondeterminism in the presence of a diverse or unknown future. In Fedor V. Fomin, Rusins Freivalds, Marta Z. Kwiatkowska, and David Peleg, editors, *ICALP 2013 (Part II)*, volume 7966 of *LNCS*, pages 89–100. Springer, 2013. doi:10.1007/978-3-642-39212-2_11.
- [BL19] Udi Boker and Karoliina Lehtinen. Good for games automata: From nondeterminism to alternation. In Wan Fokkink and Rob van Glabbeek, editors, *CONCUR 2019*, volume 140 of *LIPIcs*, pages 19:1–19:16. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2019. URL: <http://drops.dagstuhl.de/opus/volltexte/2019/10921>, doi:10.4230/LIPIcs.CONCUR.2019.19.

- [BL21] Udi Boker and Karoliina Lehtinen. History determinism vs. good for gameness in quantitative automata. In *FSTTCS 2021*, volume 213 of *LIPICs*, pages 35:1–35:20. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2021.
- [Bü64] J. Richard Büchi. Regular canonical systems. *Archiv für mathematische Logik und Grundlagenforschung*, 6(3):91–111, Apr 1964. doi:10.1007/BF01969548.
- [CG77a] Rina S. Cohen and Arie Y. Gold. Theory of omega-languages. I. Characterizations of omega-context-free languages. *J. Comput. Syst. Sci.*, 15(2):169–184, 1977. doi:10.1016/S0022-0000(77)80004-4.
- [CG77b] Rina S. Cohen and Arie Y. Gold. Theory of omega-languages. II. A study of various models of omega-type generation and recognition. *J. Comput. Syst. Sci.*, 15(2):185–208, 1977. doi:10.1016/S0022-0000(77)80005-6.
- [CG78] Rina S. Cohen and Arie Y. Gold. Omega-computations on deterministic pushdown machines. *J. Comput. Syst. Sci.*, 16(3):275–300, 1978. doi:10.1016/0022-0000(78)90019-3.
- [CH14] Arnaud Carayol and Matthew Hague. Saturation algorithms for model-checking pushdown systems. In Zoltán Ésik and Zoltán Fülöp, editors, *AFL 2014*, volume 151 of *EPTCS*, pages 1–24, 2014. doi:10.4204/EPTCS.151.1.
- [Col09] Thomas Colcombet. The theory of stabilisation monoids and regular cost functions. In Susanne Albers, Alberto Marchetti-Spaccamela, Yossi Matias, Sotiris E. Nikolettseas, and Wolfgang Thomas, editors, *ICALP 2009, (Part II)*, volume 5556 of *LNCS*, pages 139–150. Springer, 2009. doi:10.1007/978-3-642-02930-1_12.
- [Col13] Thomas Colcombet. *Fonctions régulières de coût*. Habilitation, École Doctorale de Sciences Mathématiques de Paris Centre, 2013.
- [DDK20] Manfred Droste, Sven Dziadek, and Werner Kuich. Logic for ω -pushdown automata. *Inf. Comput.*, 2020. Article in press. doi:https://doi.org/10.1016/j.ic.2020.104659.
- [Fin01] Olivier Finkel. Topological properties of omega context-free languages. *Theor. Comput. Sci.*, 262(1):669–697, 2001. doi:10.1016/S0304-3975(00)00405-9.
- [Fri10] Wladimir Fridman. Formats of winning strategies for six types of pushdown games. In Angelo Montanari, Margherita Napoli, and Mimmo Parente, editors, *GANDALF 2010*, volume 25 of *EPTCS*, pages 132–145, 2010. doi:10.4204/EPTCS.25.14.
- [GJLZ21] Shibashis Guha, Ismäl Jecker, Karoliina Lehtinen, and Martin Zimmermann. A bit of nondeterminism makes pushdown automata expressive and succinct. *arXiv*, 2105.02611, 2021.
- [GS53] David Gale and F. M. Stewart. Infinite games with perfect information. In Harold William Kuhn and Albert William Tucker, editors, *Contributions to the Theory of Games (AM-28), Volume II*, chapter 13, pages 245–266. Princeton University Press, 1953.
- [GTW02] Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *LNCS*. Springer, 2002. doi:10.1007/3-540-36387-4.
- [HMU07] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to automata theory, languages, and computation, 3rd Edition*. Pearson international edition. Addison-Wesley, 2007.
- [HP06] Thomas A. Henzinger and Nir Piterman. Solving games without determinization. In Zoltán Ésik, editor, *CSL 2006*, volume 4207 of *LNCS*, pages 395–410. Springer, 2006. doi:10.1007/11874683_26.
- [KM15] Denis Kuperberg and Michal Skrzypczak. On determinisation of good-for-games automata. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *ICALP 2015 (Part II)*, volume 9135 of *LNCS*, pages 299–310. Springer, 2015. doi:10.1007/978-3-662-47666-6_24.
- [LMS04] Christof Löding, P. Madhusudan, and Olivier Serre. Visibly pushdown games. In Kamal Lodaya and Meena Mahajan, editors, *FSTTCS 2004*, volume 3328 of *LNCS*, pages 408–420. Springer, 2004. doi:10.1007/978-3-540-30538-5_34.
- [Löd15] Christof Löding. Simplification problems for deterministic pushdown automata on infinite words. *Int. J. Found. Comput. Sci.*, 26(08):1041–1068, 2015.
- [LR12] Christof Löding and Stefan Repke. Regularity problems for weak pushdown ω -automata and games. In Branislav Rován, Vladimiro Sassone, and Peter Widmayer, editors, *MFCS 2012*, volume 7464 of *LNCS*, pages 764–776. Springer, 2012. doi:10.1007/978-3-642-32589-2_66.

- [LST94] Clemens Lautemann, Thomas Schwentick, and Denis Thérien. Logics for context-free languages. In Leszek Pacholski and Jerzy Tiuryn, editors, *CSL 1994*, volume 933 of *LNCS*, pages 205–216. Springer, 1994. doi:10.1007/BFb0022257.
- [LZ20] Karoliina Lehtinen and Martin Zimmermann. Good-for-games ω -pushdown automata. In Holger Hermanns, Lijun Zhang, Naoki Kobayashi, and Dale Miller, editors, *LICS 2020*, pages 689–702. ACM, 2020. doi:10.1145/3373718.3394737.
- [Val73] Leslie Valiant. *Decision procedures for families of deterministic pushdown automata*. PhD thesis, University of Warwick, 1973.
- [Wal01] Igor Walukiewicz. Pushdown processes: Games and model-checking. *Inf. Comput.*, 164(2):234–263, 2001. doi:10.1006/inco.2000.2894.

APPENDIX A. ANOTHER CONTEXTFREE LANGUAGE THAT IS NOT GOOD-FOR-GAMES

Let $h: \{0, 1, \#\}^* \rightarrow \{0, 1\}^*$ be the homomorphism induced by $h(0) = 0$, $h(1) = 1$, and $h(\#) = \varepsilon$. Define

$$P = \{v\#^\omega \mid h(v) = xx^R \text{ for some } x \in \{0, 1\}^*\},$$

where x^R denotes the reversal of x . It is straightforward to construct an ω -PDA recognizing P , thereby showing $P \in \omega\text{-CFL}$. We show $P \notin \omega\text{-GFG-CFL}$.

Towards a contradiction, let us assume that P is recognized by an $\omega\text{-GFG-PDA } \mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, \Omega)$ with resolver $r: \Delta^* \times \Sigma \rightarrow \Delta$. Note that r induces a unique run ρ_w for every $w \in P$. Furthermore, if two words $w, w' \in P$ share a prefix, then their runs $\rho_w = c_0\tau_0c_1\tau_1c_2\tau_2 \cdots$ and $\rho_{w'} = c'_0\tau'_0c'_1\tau'_1c'_2\tau'_2 \cdots$ also share a prefix. More formally, if

$$w(0) \cdots w(n) = w'(0) \cdots w'(n)$$

for some $n \geq 0$, then we have $\tau_0 \cdots \tau_k = \tau'_0 \cdots \tau'_k$ and thus $c_0 \cdots c_{k+1} = c'_0 \cdots c'_{k+1}$ for every k such that $\ell(\tau_0 \cdots \tau_k) = w(0) \cdots w(n-1)$. Here, the -1 stems from the fact that r resolves nondeterminism based on the next input letter to be processed.

Let $\rho = c_0\tau_0c_1 \cdots c_k\tau_kc_{k+1}$ be an alternating sequence of configurations and transitions, and let $a \in \{0, 1, \#\}$. We say that (ρ, a) is r -consistent, if

$$\tau_n = r(\tau_0 \cdots \tau_{n-1}, v(|\ell(\tau_0 \cdots \tau_{n-1})|))$$

and c_n is the configuration reached by \mathcal{P} after the transitions $\tau_0 \cdots \tau_{n-1}$, where $v = \ell(\rho)a$. Now, an r -consistent pair (ρ, a) has Property M if the following holds for every $w \in \{0, 1, \#\}^*\#^\omega$: Let $c'_0\tau'_0c'_1\tau'_1c'_2\tau'_2 \cdots$ be the unique run of \mathcal{P} on $\ell(\rho)aw$ induced by r . Then, we require $\text{sh}(c'_{k'+1}) \geq \text{sh}(c_{k+1})$ for all $k' > k$, i.e., that $k+1$ is a step (see Page 9). Note that we have $\tau'_0 \cdots \tau'_k = \tau_0 \cdots \tau_k$ and $c_0 \cdots c_{k+1} = c'_0 \cdots c'_{k+1}$ by definition.

Lifting Property M to words, we say that $v(0) \cdots v(n) \in \{0, 1, \#\}^+$ has property M if there is a run prefix ρ with $\ell(\rho) = v(0) \cdots v(n-1)$ such that $(\rho, v(n))$ is r -consistent and has Property M.

We show that for every $v \in \{0, 1, \#\}^+$ there is a $v' \in \{0, 1, \#\}^*$ such that vv' has Property M. To this end, assume $v(0) \cdots v(n)$ does not have Property M. Let ρ be a run prefix of \mathcal{P} on $v(0) \cdots v(n-1)$ so that $(\rho, v(n))$ is r -consistent. As v is a prefix of some word in P , such a ρ exists. By definition, (ρ, a) does not have property M, i.e., there is a run induced by r processing a prolongation in $v\{0, 1, \#\}^*\#^\omega$ that reaches a stack height strictly smaller than the stack height of the last configuration of ρ . As stack heights are bounded from below, there is a minimal stack height that is assumed by such runs. Let vv' with $v' \in \{0, 1, \#\}^*$ be a word processed by such a run prefix ρ' ending with a minimal stack

height along all runs considered. Then, $vv'0$ has Property M, as $(\rho', 0)$ has property M, as after ρ' no strictly smaller stack height is reached by runs induced by r .

Now, fix $n = 3|Q|(|\Gamma + 1|) + 1$ and define $v_j = 01^j0$ for every j in the range $1 \leq j \leq n$. As shown above, for every such j , there is a v'_j such that $v_jv'_j$ has Property M.

Now, each $v_jv'_j(v_jv'_j)^R\#\omega$ is in L , i.e., the unique run

$$c_0^j\tau_0^j c_1^j\tau_1^j c_2^j\tau_2^j \dots$$

of \mathcal{P} on $v_jv'_j(v_jv'_j)^R\#\omega$ induced by r is accepting. As each $v_jv'_j$ has Property M, there are prefixes $c_0^j\tau_0^j \dots c_{k_j}^j\tau_{k_j}^j c_{k_j+1}^j$ of $c_0^j\tau_0^j c_1^j\tau_1^j c_2^j\tau_2^j \dots$ processing $v_jv'_j$ without its last letter a_j such that $(c_0^j\tau_0^j \dots c_{k_j}^j\tau_{k_j}^j c_{k_j+1}^j, a_j)$ has Property M.

Now, there are $j_0 \neq j_1$ such that the configurations $c_{k_{j_0}+1}^{j_0}$ and $c_{k_{j_1}+1}^{j_1}$ coincide on their state from Q and their top stack symbol from Γ and such that $a_{j_0} = a_{j_1}$.

Consider the sequence

$$\rho^* = \tau_0^{j_0} \dots \tau_{k_{j_0}}^{j_0} \tau_{k_{j_1}+1}^{j_1} \tau_{k_{j_1}+2}^{j_1} \tau_{k_{j_1}+3}^{j_1} \dots$$

We claim that ρ^* induces an accepting run of \mathcal{P} on $\bar{w} = v_{j_0}v'_{j_0}(v_{j_1}v'_{j_1})^R\#\omega$. We have

$$v_{j_0}v'_{j_0}(v_{j_1}v'_{j_1})^R = 01^{j_0}0v'_{j_0}(v_{j_1})^R01^{j_1}0,$$

which is not of the form vv^R after removing $\#$'s. Hence, this step completes the proof, as $\bar{w} \notin P$ is accepted by \mathcal{P} , yielding the desired contradiction.

The sequence ρ^* satisfies the acceptance condition, as it shares a suffix with the sequence of transitions of the accepting run $c_0^{j_0}\tau_0^{j_0} c_1^{j_1}\tau_1^{j_1} c_2^{j_1}\tau_2^{j_1} \dots$. Furthermore, we have

$$\ell(\rho^*) = \ell(\tau_0^{j_0} \dots \tau_{k_{j_0}}^{j_0}) \ell(\tau_{k_{j_1}+1}^{j_1} \tau_{k_{j_1}+2}^{j_1} \tau_{k_{j_1}+3}^{j_1} \dots)$$

where $\ell(\tau_0^{j_0} \dots \tau_{k_{j_0}}^{j_0})$ is equal to $v_{j_0}v'_{j_0}$ without its last letter a_{j_0} , and where

$$\ell(\tau_{k_{j_1}+1}^{j_1} \tau_{k_{j_1}+2}^{j_1} \tau_{k_{j_1}+3}^{j_1} \dots) = a_{j_1}(v_{j_1}v'_{j_1})^R\#\omega.$$

The concatenation of these two words is indeed \bar{w} , as we have $a_{j_0} = a_{j_1}$ by construction.

Finally, as $(c_0^{j_1}\tau_0^{j_1} \dots c_{k_{j_1}}^{j_1}\tau_{k_{j_1}}^{j_1} c_{k_{j_1}+1}^{j_1}, a_{j_1})$ has Property M, the run $c_0^{j_1}\tau_0^{j_1} c_1^{j_1}\tau_1^{j_1} c_2^{j_1}\tau_2^{j_1} \dots$ does after position k_{j_1} not depend on the complete stack content at that position, but only on the state and the top stack symbol of $c_{k_{j_1}+1}^{j_1}$. Hence, appending the run suffix resulting from applying the transitions

$$\tau_{k_{j_1}+1}^{j_1} \tau_{k_{j_1}+2}^{j_1} \tau_{k_{j_1}+3}^{j_1} \dots$$

after $(c_0^{j_0}\tau_0^{j_0} \dots c_{k_{j_0}}^{j_0}\tau_{k_{j_0}}^{j_0} c_{k_{j_0}+1}^{j_0}, a_{j_0})$ (recall that $c_{k_{j_0}+1}^{j_0}$ and $c_{k_{j_1}+1}^{j_1}$ have the same state and top stack symbol) yields indeed a run of \mathcal{P} (but not necessarily induced by r).

APPENDIX B. SYNTHESIS FROM ω -GFG-CFL SPECIFICATIONS

In Section 4, we showed that the winner of a Gale-Stewart game with an ω -GFG-CFL winning condition can be determined in exponential time. Here, we consider the synthesis problem: If Player 2 wins a Gale-Stewart game with an ω -GFG-CFL winning condition, compute a (finitely represented) winning strategy for her. We show that this problem can be solved in exponential time as well, as for Gale-Stewart games with ω -DCFL winning conditions.

Assume we are given an ω -GFG-PDA \mathcal{P} such that Player 2 wins $\mathcal{G}(L(\mathcal{P}))$. In the proof of Theorem 4.2, we have constructed a polynomially-sized ω -DPDA \mathcal{P}_d such that Player 2 wins $\mathcal{G}(L(\mathcal{P}))$ if and only if she wins $\mathcal{G}(L(\mathcal{P}_d))$. We show how a finitely represented winning strategy for Player 2 in $\mathcal{G}(L(\mathcal{P}_d))$, which exists if she wins $\mathcal{G}(L(\mathcal{P}_d))$, can be turned into the desired strategy for $\mathcal{G}(L(\mathcal{P}))$.

In our context, we need a slight adaption of the pushdown transducer model we introduced in Section 7 to implement resolvers. Here, a pushdown transducer (PDT) is a tuple $\mathcal{T} = (\bar{Q}, \Sigma_I, \bar{\Gamma}, \bar{q}_I, \bar{\Delta}, \Sigma_O, \bar{\lambda})$ such that $(\bar{Q}, \Sigma_I, \bar{\Gamma}, \bar{q}_I, \bar{\Delta})$ is a deterministic PDA (without coloring) over Σ_I , Σ_O is an output alphabet, and $\bar{\lambda}: \bar{Q} \rightarrow \Sigma_O$ is a (partial) output function. Note that we only change the signature of the output function. Hence, runs of \mathcal{T} are defined as before. In particular, we still require them to be maximal, i.e., they end in configurations without enabled ε -transitions. Thus, we say that \mathcal{T} computes the (partial) function $\sigma_{\mathcal{T}}: \Sigma_I^* \rightarrow \Sigma_O$ given by $\sigma_{\mathcal{T}}(w) = \bar{\lambda}(q)$, where q is the state of the configuration the run of \mathcal{T} on w ends in.

In a Gale-Stewart game $\mathcal{G}(L)$ with $L \subseteq (\Sigma_1 \times \Sigma_2)^\omega$, a strategy for Player 2 is a function from Σ_1^+ to Σ_2 . Such strategies can be implemented by a pushdown transducer with input alphabet $\Sigma_I = \Sigma_1$ and output alphabet $\Sigma_O = \Sigma_2$.

Proposition B.1 [Fri10, Wal01]. *Let \mathcal{P}_d be an ω -DPDA. If Player 2 wins $\mathcal{G}(L(\mathcal{P}_d))$, then she has a winning strategy that is implemented by a PDT. Furthermore, such a PDT can be computed in exponential time in $|\mathcal{P}_d|$.*

Now, we are ready to prove Theorem 4.4.

Proof. Given an ω -GFG-PDA \mathcal{P} , let the ω -DPDA \mathcal{P}_d be defined as in the proof of Theorem 4.2. Then, the game $\mathcal{G}(L(\mathcal{P}_d))$ is won by Player 2 and she has a winning strategy for $\mathcal{G}(\mathcal{P}_d)$ that is implemented by a PDT, which can be computed in exponential time (in the size of \mathcal{P} , as \mathcal{P}_d is only polynomially larger than \mathcal{P}). Now, we show that the transformation $\sigma \mapsto \sigma_{-d}$ presented in the proof of Theorem 4.2 is effective, if σ is implemented by a PDT. That is, we turn a PDT \mathcal{T} implementing σ into a PDT \mathcal{T}_{-d} implementing σ_{-d} .

So, consider a PDT $\mathcal{T} = (\bar{Q}, (\Sigma_1)_d, \bar{\Gamma}, \bar{q}_I, \bar{\Delta}, (\Sigma_2)_d, \bar{\lambda})$ implementing a winning strategy σ for Player 2 in $\mathcal{G}(\mathcal{P}_d)$ where we write $(\Sigma_1)_d$ for Σ_1 and write $(\Sigma_2)_d$ for $\Sigma_2 \cup \Delta$ to distinguish the alphabets of \mathcal{P} and \mathcal{P}_d . The mode of a configuration $(q, \gamma X)$ of \mathcal{T} is the pair (q, X) (note that X might be \perp). Whether a transition is enabled in a configuration only depends on its mode. Thus, we are justified to say that a mode enables a transition or not. Now, a mode is a reading mode if it enables some non- ε transition. As \mathcal{T} is deterministic, no reading mode has an enabled ε -transition. Furthermore, due to the maximality requirement on runs, every run ends in a reading mode. Finally, our construction relies on the following simple fact: If we remove all but one transition enabled by a reading mode, and turn the remaining transition into an ε -transition, then the resulting PDT is still deterministic.

First, we transform \mathcal{T} into a PDT computing the same strategy as \mathcal{T} while keeping track of its own mode using its states. This is useful as it allows us to base the output function, which formally only depends on a state of a configuration, on the current mode. Thus, the new PDT has to keep track of the topmost stack symbol, which is straightforward during transitions pushing a symbol on the stack and during transitions replacing the topmost stack symbol, but requires an additional transition to read the newly revealed stack symbol after the topmost stack symbol is popped off the stack.

Formally, define $\mathcal{T}' = (\overline{Q} \times \overline{\Gamma}_\perp \cup Q, (\Sigma_1)_d, \overline{\Gamma}, (\overline{q_I}, \perp), \overline{\Delta}', (\Sigma_2)_d, \overline{\lambda}')$ with

$$\begin{aligned} \overline{\Delta}' = & \{((q, X), X, a, (q', X''), X'X'') \mid (q, X, a, q', X'X'') \in \overline{\Delta}\} \cup \\ & \{((q, X), X, a, (q', X'), X') \mid (q, X, a, q', X') \in \overline{\Delta}\} \cup \\ & \{((q, X), X, a, q', \varepsilon) \mid (q, X, a, q', \varepsilon) \in \overline{\Delta}\} \cup \\ & \{(q, X, \varepsilon, (q, X), X) \mid q \in \overline{Q}, X \in \overline{\Gamma}_\perp\}, \end{aligned}$$

which is still deterministic. Further, \mathcal{T}' has a run on some input v if and only if \mathcal{T} has a run on v . Finally, let the run of \mathcal{T} on v end in some configuration $(q, \gamma X)$. Then, the run of \mathcal{T}' on v ends in the configuration $((q, X), \gamma X)$. Hence, both PDT's indeed compute the same function when defining $\overline{\lambda}'(q, X) = \overline{\lambda}(q)$.

Now, we define $\mathcal{T}_{-d} = (Q_{-d}, \Sigma_1, \overline{\Gamma}, (\overline{q_I}, \perp, \iota), \overline{\Delta}_{-d}, \Sigma_2, \overline{\lambda}_{-d})$ with state set

$$Q_{-d} = (\overline{Q} \times \overline{\Gamma}_\perp \cup Q) \times (\{\iota, \omega\} \cup \Sigma_2).$$

Next, we define the transition relation $\overline{\Delta}_{-d}$, whose definition refers to the fixed dummy letter $\perp \in \Sigma_1$ used in the definition of σ_{-d} (see Page 4). Intuitively, \mathcal{T}_{-d} works in three distinct phases as follows:

- The initialization phase spans the initial sequence of ε -transitions executed by \mathcal{T}' and is left for the waiting phase as soon as the first letter from Σ_1 is processed. Formally,
 - $((q, \iota), X, \varepsilon, (q', \iota), \gamma) \in \overline{\Delta}_{-d}$ for every $(q, X, \varepsilon, q', \gamma) \in \overline{\Delta}'$, and
 - $((q, \iota), X, a, (q', \omega), \gamma) \in \overline{\Delta}_{-d}$ for every $(q, X, a, q', \gamma) \in \overline{\Delta}'$ with $a \in \Sigma_1$.
- In the waiting phase, ε -transitions are simulated until none are enabled any more. At that point, \mathcal{P}' yields an output and processes some input from Σ_1 . In \mathcal{T}_{-d} , the output letter is stored and the unique transition labeled by \perp is turned into an ε -transition, which starts the delay phase. As we remove all other a -transitions with $a \neq \perp$, this preserves determinism. Formally,
 - $((q, \omega), X, \varepsilon, (q', \omega), \gamma) \in \overline{\Delta}_{-d}$ for every $(q, X, \varepsilon, q', \gamma) \in \overline{\Delta}'$ and
 - $((q, \omega), X, \varepsilon, (q', \overline{\lambda}'(q)), \gamma) \in \overline{\Delta}_{-d}$ for every $(q, X, \perp, q', \gamma) \in \overline{\Delta}'$.
- During the delay phase, ε - and \perp -transitions of \mathcal{T}' are simulated (deterministically, and without processing an input letter) until a reading mode is reached at which \mathcal{T}' outputs a non- ε -transition, signifying the end of a block. At that point, the output letter stored during the delay phase is finally output and the run returns to the waiting phase. As we remove all other transitions when turning \perp -transitions into ε -transitions, we again preserve determinism. Formally,
 - $((q, a_2), X, \varepsilon, (q', a_2), \gamma) \in \overline{\Delta}_{-d}$ for every $a_2 \in \Sigma_2$ and every $(q, X, \varepsilon, q', \gamma) \in \overline{\Delta}'$,
 - $((q, a_2), X, \varepsilon, (q', a_2), \gamma) \in \overline{\Delta}_{-d}$ for every $a_2 \in \Sigma_2$ and every $(q, X, \perp, q', \gamma) \in \overline{\Delta}'$ such that q is a reading mode of \mathcal{T} and $\overline{\lambda}'(q)$ is an ε -transition of \mathcal{P} , and
 - $((q, a_2), X, a_1, (q', \omega), \gamma) \in \overline{\Delta}_{-d}$ for every $a_2 \in \Sigma_2$ and every $(q, X, a_1, q', \gamma) \in \overline{\Delta}'$ such that q is a reading mode of \mathcal{T} , but $\overline{\lambda}'(q)$ is a non- ε -transition of \mathcal{P} .

Finally, we define $\overline{\lambda}_{-d}(q, a_2) = a_2$ for every state q of \mathcal{T}' that is a reading mode of \mathcal{T} , but $\overline{\lambda}'(q)$ is a non- ε -transition of \mathcal{P} . Thus, we indeed output the stored letter at the end of the delay phase as described above.

Recall that in the proof of Theorem 4.2, we have turned the winning strategy σ for Player 2 in $\mathcal{G}(L(\mathcal{P}_d))$ into a winning strategy σ_{-d} for her in $\mathcal{G}(L(\mathcal{P}))$. To this end, we first defined a function mapping $v \in \Sigma_1^+$ to some $v_d \in (\Sigma_1')^+$ as follows: $a_d = a$ for $a \in \Sigma_1$

and $(va)_d = v_d \cdot^n a$ for some n that is independent of a . With these definitions, we defined $\sigma_{-d}(v) = \sigma(v_d)$. In the remainder of the proof, we show that \mathcal{T}_{-d} indeed computes σ_{-d} .

So, consider some $v \in \Sigma_1^+$ and the unique n such that $(va)_d = v_d \cdot^n a$ for all $a \in \Sigma_1$. Then, we have $\sigma(v_d) = \bar{\lambda}'(q)$, where q is the state of the last configuration of the run ρ of \mathcal{T}' on v_d . Furthermore, consider the run ρ' of \mathcal{T}' on $v_d \cdot^n$, which is a proper extension of ρ ending in some state q' . By definition of n , q' is a reading mode of \mathcal{T} and $\bar{\lambda}'(q')$ is a non- ε -transition of \mathcal{P} .

The run ρ' can be turned into the unique run of \mathcal{T}_{-d} on v by switching between the phases appropriately:

- The initialization phase ends with the transition processing the first letter of v_d , which starts a waiting phase.
- Each waiting phase lasts until a state is reached at which a letter in Σ_2 is output by \mathcal{T}' . This letter is stored and a delay phase is started.
- Each delay phase lasts until a state is reached at which a non- ε -transition is output by \mathcal{T}' .

In particular, after the last letter of v is processed, a waiting phase starts, which is ended by switching from state (q, ω) to a state storing $\bar{\lambda}'(q)$. The last delay phase stores this letter until the end of the run, at which point it yields the output $\bar{\lambda}'(q)$ of \mathcal{T}_{-d} on v . Thus, \mathcal{T}_{-d} indeed implements σ_{-d} . \square

Finally, let us consider the situation of Player 1 in games with ω -GFG-CFL winning conditions, where we define strategies and strategies implemented by PDT's for Player 1 as expected. As ω -GFG-CFL is not closed under complementation, we cannot dualize a game with a ω -GFG-CFL winning condition by swapping the roles of the players and complementing the winning condition to obtain a winning strategy for Player 1. Instead, the situation is asymmetric.

Theorem B.2. *There is a Gale-Stewart game with an ω -GFG-CFL winning condition that is won by Player 1, but not with a winning strategy implementable by a PDT.*

Proof. In Section 3, we have introduced the language L_{ss} over $I \times I$ for $I = \{+, -, 0\}$ and the words $x_1 = \binom{+}{0} \binom{+}{-}$ and $x_2 = \binom{0}{+} \binom{-}{+}$. Now, consider the game $\mathcal{G}(L)$ with

$$L = \left\{ \binom{w}{\#^\omega} \mid w \in L_{ss} \text{ or } w \notin \{x_1, x_2\}^\omega \right\},$$

which is in ω -GFG-CFL due to Theorem 5.3.

Recall that the ω -word

$$w_{ss} = x_1 (x_2)^3 (x_1)^7 (x_2)^{15} (x_1)^{31} (x_2)^{63} \dots$$

defined on Page 9 is not in L_{ss} . Thus, Player 1 wins $\mathcal{G}(L)$ by producing the word w_{ss} .

Now, consider a PDT \mathcal{T} implementing a strategy σ for Player 1 in $\mathcal{G}(L)$. Let $w \binom{w}{\#^\omega}$ be an outcome that is consistent with σ . The run ρ of \mathcal{T} generating this outcome has infinitely many steps. Now, pick two such steps, such that the configurations at these steps coincide on their state, top stack symbol, and such that at least three outputs are generated between the steps. Then, the run obtained by repeating the sequence of transitions taken between the steps generates an outcome consistent with σ . Now, this outcome has a positive energy level in one component, as every infix of length at least 3 of a word built by concatenating copies of the x_i has a strictly positive energy level in one component. Hence, σ is not a winning strategy for Player 1 in $\mathcal{G}(L)$.

Hence, no PDT implements a winning strategy for Player 1 in $\mathcal{G}(L)$. \square

Note that Player 1 *does* have a winning strategy implementable by a PDT in the game $\mathcal{G}(L(\mathcal{P}_d))$ with ω -DCFL winning condition obtained from $\mathcal{G}(L)$ (see the proof of Theorem 4.2). In the game $\mathcal{G}(L(\mathcal{P}_d))$, Player 2 has to construct a run of an ω -GFG-PDA recognizing L , which introduces enough information for Player 1 to win with a strategy implementable by a PDT. However, in $\mathcal{G}(L)$, he does not have access to that additional information, and therefore no winning strategy implementable by a PDT.