



HAL
open science

A Bit of Nondeterminism Makes Pushdown Automata Expressive and Succinct

Shibashis Guha, Ismaël Jecker, Karoliina Lehtinen, Martin Zimmermann

► **To cite this version:**

Shibashis Guha, Ismaël Jecker, Karoliina Lehtinen, Martin Zimmermann. A Bit of Nondeterminism Makes Pushdown Automata Expressive and Succinct. 46th International Symposium on Mathematical Foundations of Computer Science (MFCS 2021), Aug 2021, Tallinn, Estonia. 10.4230/LIPIcs.MFCS.2021.53 . hal-03320016

HAL Id: hal-03320016

<https://hal.science/hal-03320016>

Submitted on 2 Mar 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Bit of Nondeterminism Makes Pushdown Automata Expressive and Succinct

Shibashis Guha¹, Ismaël Jecker², Karoliina Lehtinen³ and Martin Zimmermann⁴

¹ Tata Institute of Fundamental Research, Mumbai, India

shibashis.guha@tifr.res.in

² IST Austria, Klosterneuburg, Austria

ismael.jecker@ist.ac.at

³ CNRS, Aix-Marseille University and University of Toulon, LIS, Marseille, France

lehtinen@lis-lab.fr

⁴ University of Liverpool, Liverpool, United Kingdom

{Martin.Zimmermann}@liverpool.ac.uk

Abstract. We study the expressiveness and succinctness of good-for-games pushdown automata (GFG-PDA) over finite words, that is, pushdown automata whose nondeterminism can be resolved based on the run constructed so far, but independently of the remainder of the input word.

We prove that GFG-PDA recognise more languages than deterministic PDA (DPDA) but not all context-free languages (CFL). This class is orthogonal to unambiguous CFL. We further show that GFG-PDA can be exponentially more succinct than DPDA, while PDA can be double-exponentially more succinct than GFG-PDA. We also study GFGness in visibly pushdown automata (VPA), which enjoy better closure properties than PDA, and for which we show GFGness to be EXPTIME-complete. GFG-VPA can be exponentially more succinct than deterministic VPA, while VPA can be exponentially more succinct than GFG-VPA. Both of these lower bounds are tight.

Finally, we study the complexity of resolving nondeterminism in GFG-PDA. Every GFG-PDA has a positional resolver, a function that resolves nondeterminism and that is only dependant on the current configuration. Pushdown transducers are sufficient to implement the resolvers of GFG-VPA, but not those of GFG-PDA. GFG-PDA with finite-state resolvers are determinisable.

1 Introduction

Nondeterminism adds both expressiveness and succinctness to deterministic pushdown automata. Indeed, the class of context-free languages (CFL), recognised by nondeterministic pushdown automata (PDA), is strictly larger than the class of deterministic context-free languages (DCFL), recognised by deterministic pushdown automata (DPDA), both over finite and infinite words. Even when restricted to languages in DCFL, there is no computable bound on the relative succinctness of PDA [15,39]. In other words, nondeterminism is remarkably powerful, even for representing deterministic languages. The cost of such succinct representations is algorithmic: problems such as universality and solving games with a CFL winning condition are undecidable for PDA [12,19], while they are decidable for DPDA [40]. Intermediate forms of automata that lie between deterministic and nondeterministic models have the potential to mitigate some of the disadvantages of fully nondeterministic automata while retaining some of the benefits of the deterministic ones.

Unambiguity and bounded ambiguity, for example, restrict nondeterminism by requiring words to have at most one or at most k , for some fixed k , accepting runs. Holzer and Kutrib survey the noncomputable succinctness gaps between unambiguous PDA and both PDA and DPDA [18], while Okhotin and Salomaa show that unambiguous visibly pushdown automata are exponentially more succinct than DPDA [32]. Universality of unambiguous PDA is decidable, as it is decidable for unambiguous context-free grammars [34], which are effectively equivalent [17]. However, to the best of our knowledge, unambiguity is not known to reduce the algorithmic complexity of solving games with a context-free winning condition.

Another important type of restricted nondeterminism that is known to reduce the complexity of universality and solving games has been studied under the names of good-for-games (GFG) nondeterminism [16] and history-determinism [11]. Intuitively, a nondeterministic automaton is GFG if its nondeterminism can be resolved on-the-fly, i.e. without knowledge of the remainder of the input word to be processed.

For finite automata on finite words, where nondeterminism adds succinctness, but not expressiveness, GFG nondeterminism does not even add succinctness: every GFG-NFA contains an equivalent DFA [6], which can be obtained by pruning transitions from the GFG-NFA. Thus, GFG-NFA cannot be more succinct than DFA. But for finite automata on infinite words, where nondeterminism again only adds succinctness, but not expressiveness, GFG coBüchi automata can be exponentially more succinct than deterministic automata [23]. Finally, for certain quantitative automata over infinite words, GFG nondeterminism adds as much expressiveness as arbitrary nondeterminism [11].

Recently, pushdown automata on infinite words with GFG nondeterminism (ω -GFG-PDA) were shown to be strictly more expressive than ω -DPDA, while universality and solving games for ω -GFG-PDA are not harder than for ω -DPDA [25]. Thus, GFG nondeterminism adds expressiveness without increasing the complexity of these problems, i.e. pushdown automata with GFG nondeterminism induce a novel and intriguing class of context-free ω -languages.

Here, we continue this work by studying the expressiveness *and* succinctness of PDA over finite words. While the decidability results for ω -GFG-PDA on infinite words also hold for GFG-PDA on finite words, the separation argument between ω -GFG-PDA and ω -DPDA depends crucially on combining GFG nondeterminism with the coBüchi acceptance condition. Since this condition is only relevant for infinite words, the separation result does not transfer to the setting of finite words.

Nevertheless, we prove that GFG-PDA are more expressive than DPDA, yielding the first class of automata on finite words where GFG nondeterminism adds expressiveness. The language witnessing the separation is remarkably simple, in contrast to the relatively subtle argument for the infinitary result [25]: the language $\{a^i b^j c^k \mid k \leq \max(i, j)\}$ is recognised by a GFG-PDA but not by a DPDA. This yields a new class of languages, those recognised by GFG-PDA over finite words, for which universality and solving games are decidable. We also show that this class is incomparable with unambiguous context-free languages.

We then turn our attention to succinctness of GFG-PDA. We show that the succinctness gap between DPDA and GFG-PDA is at least exponential, while the gap between GFG-PDA and PDA is at least double-exponential. These results hold already for finite words.

To the best of our knowledge, both our expressiveness and our succinctness results are the first examples of good-for-games nondeterminism being used effectively over finite, rather than infinite, words (recall that all GFG-NFA are determinisable by pruning). Also, this is the first succinctness result for good-for-games automata that does not depend on the infinitary coBüchi acceptance condition, which was used to show the exponential succinctness of GFG coBüchi automata, as compared to deterministic ones [23].

We then study an important subclass of GFG-PDA, namely, GFG visibly pushdown automata (VPA), in which the stack behaviour (push, pop, skip) is determined by the input letter only. GFG-VPA enjoy good closure properties: they are closed under complement, union and intersection. We show that there is an exponential succinctness gap between deterministic VPA and GFG-VPA, as well as between GFG-VPA and VPA. Both of these are tight, as VPA, and therefore GFG-VPA as well, admit an exponential determinisation procedure [2]. Furthermore, we show that GFGness of VPA is decidable in EXPTIME. This makes GFG-VPA a particularly interesting class of PDA as they are recognisable, succinct, have good closure properties and deciding universality and solving games are both in EXPTIME. In contrast, solving ω -VPA games is 2EXPTIME-complete [27]. We also relate the problem of checking GFGness with the *good-enough synthesis* [1] or *uniformization* problem [9], which we show to be EXPTIME-complete for DVPA and GFG-PDA.

Nondeterminism in GFG automata is resolved on-the-fly, i.e. the next transition to be taken only depends on the run prefix constructed so far and the next letter to be processed. Thus, the complexity of a resolver, mapping run prefixes and letters to transitions, is a natural complexity measure for GFG automata. For example, finite GFG automata (on finite and infinite words) have a finite-state resolver [16]. For pushdown automata with their infinite configuration space, the situation is markedly different: On one hand, we show that GFG-PDA admit positional resolvers, that is, resolvers that depend only on the current configuration, rather than on the entire run prefix produced so far. Note that this result only holds for GFG-PDA over finite words, but not for ω -GFG-PDA. Yet, positionality does not imply that resolvers are simple to implement. We show that there are GFG-PDA that do not admit a resolver implementable by a pushdown transducer. In contrast, all GFG-VPA admit pushdown resolvers, again showing that GFG-VPA are better behaved than general GFG-PDA. Finally, GFG-PDA with finite-state resolvers are determinisable.

All proofs omitted due to space restrictions can be found in the appendix.

Related work The notion of GFG nondeterminism has emerged independently several times, at least as Colcombet’s history-determinism [11], in Piterman and Henzinger’s GFG automata [16], and as Kupferman, Safra, and Vardi’s nondeterminism for recognising derived languages, that is, the language of trees of which all branches are in a regular language [24]. Related notions have also emerged in the context of XML document parsing. Indeed, preorder typed visibly pushdown languages and 1-pass preorder typeable tree languages, considered by Kumar, Madhusudan, and Viswanathan [21] and Martens, Neven, Schwentick, and Bex [28] respectively, also consider nondeterminism which can be resolved on-the-fly. However, the restrictions there are stronger than simple GFG nondeterminism, as they also require the typing to be unique, roughly corresponding to unambiguity in automata models and grammars. This motivates the further study of unambiguous GFG automata, although this remains out of scope for the present paper. The XML extension AXML has also inspired Active Context Free Games [30], in which one player, aiming to produce a word within a target regular language, chooses positions on a word and the other player chooses a rewriting rule from a context-free grammar. Restricting the strategies of the first player to moving from left to right makes finding the winner decidable [30,5]; however, since the player still knows the future of the word, this restriction is not directly comparable to GFG nondeterminism.

Unambiguity, or bounded ambiguity, is an orthogonal way of restricting nondeterminism by limiting the number of permitted accepting runs per word. For regular languages, it leads to polynomial equivalence and containment algorithms [38]. Minimization remains NP-complete for both unambiguous automata [20,4] and GFG automata [36] (at least when acceptance is defined on states, see [33]). On pushdown automata, increasing the permitted degree of ambiguity leads to both greater expressiveness and unbounded succinctness [17]. Finally, let us mention two more ways of measuring—and restricting—nondeterminism in PDA: bounded nondeterminism, as studied by Herzog [17] counts the branching in the run-tree of a word, while the minmax measure [35,14] counts the number of nondeterministic guesses required to accept a word. The natural generalisation of GFGness as the *width* of an automaton [22] has not yet, to the best of our knowledge, been studied for PDA.

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^+ . The empty word is denoted by ε , the length of a word w is denoted by $|w|$, and the n^{th} letter of w is denoted by $w(n)$ (starting with $n = 0$). A language over Σ is a subset of Σ^* .

For alphabets Σ_1, Σ_2 , we extend functions $f: \Sigma_1 \rightarrow \Sigma_2^*$ homomorphically to words over Σ_1 via $f(w) = f(w(0))f(w(1))f(w(2))\dots$.

2.1 Pushdown automata

A pushdown automaton (PDA for short) $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ 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 set F of final states. 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 sequence ρ over Δ , the word $\ell(\rho) \in \Sigma^*$ is defined by applying ℓ homomorphically to every transition. We take the size of \mathcal{P} to be $|Q| + |\Gamma|$.¹

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 initial configuration is (q_I, \perp) .

The set of modes of \mathcal{P} is $Q \times \Gamma_\perp$. A mode (q, X) enables all transitions of the form (q, X, a, q', γ') for some $a \in \Sigma_\varepsilon$, $q' \in Q$, and $\gamma' \in \Gamma_\perp^{\leq 2}$. The mode of a configuration $c = (q, \gamma X)$ is (q, X) . A transition τ

¹ Note that we prove exponential succinctness gaps, so the exact definition of the size is irrelevant, as long as it is polynomial in $|Q|$ and $|\Gamma|$. Here, we pick the sum for the sake of simplicity.

is enabled by c if it is enabled by c 's mode. In this case, we write $(q, \gamma X) \xrightarrow{\tau} (q', \gamma \gamma')$, where $\tau = (q, X, a, q', \gamma')$.

A run of \mathcal{P} is a finite sequence $\rho = c_0 \tau_0 c_1 \tau_1 \cdots c_{n-1} \tau_{n-1} c_n$ of configurations and transitions with c_0 being the initial configuration and $c_{n'} \xrightarrow{\tau_{n'}} c_{n'+1}$ for every $n' < n$. The run ρ is a run of \mathcal{P} on $w \in \Sigma^*$, if $w = \ell(\rho)$. We say that ρ is accepting if it ends in a configuration whose state is final. The language $L(\mathcal{P})$ recognized by \mathcal{P} contains all $w \in \Sigma^*$ such that \mathcal{P} has an accepting run on w .

Remark 1. Let $c_0 \tau_0 c_1 \tau_1 \cdots c_{n-1} \tau_{n-1} c_n$ be a run of \mathcal{P} . Then, the sequence $c_0 c_1 \cdots c_{n-1} c_n$ of configurations is uniquely determined by the sequence $\tau_0 \tau_1 \cdots \tau_{n-1}$ 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 $\tau_{n'}$ is not enabled in $c_{n'}$).

We say that a PDA \mathcal{P} is deterministic (DPDA) if

- every mode (q, X) of \mathcal{P} enables at most one a -transition for every $a \in \Sigma \cup \{\varepsilon\}$, and
- for every mode (q, X) of \mathcal{P} , if it enables some ε -transition, then it does not enable any Σ -transition.

Hence, for every input and for every run prefix on it there is a unique enabled transition to continue the run. Still, due to the existence of ε -transitions, a DPDA can have more than one run on a given input. However, these only differ by trailing ε -transitions.

The class of languages recognized by PDA is denoted by CFL, the class of languages recognized by DPDA by DCFL.

Example 1. The PDA \mathcal{P} depicted in Figure 1 recognizes the language $\{ac^n d^n a \mid n \geq 1\} \cup \{bc^n d^{2n} b \mid n \geq 1\}$. Note that while \mathcal{P} is nondeterministic, $L(\mathcal{P})$ is in DCFL.




Fig. 1. The PDA \mathcal{P} from Example 1. Grey states are final, and X is an arbitrary stack symbol.

2.2 Good-for-games Pushdown Automata

Here, we introduce good-for-games pushdown automata on finite words (GFG-PDA for short), nondeterministic pushdown automata whose nondeterminism can be resolved based on the run prefix constructed so 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 PDA \mathcal{P} from Example 1. It is nondeterministic, but knowing whether the first transition of the run processed an a or a b allows the nondeterminism to be resolved in a configuration of the form $(q, \gamma N)$ when processing a d : in the former case, take the transition to state q_1 , in the latter case the transition to state q_2 . 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.

Formally, a PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ 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 $w \in L(\mathcal{P})$, there is an accepting run $\rho = c_0 \tau_0 \cdots \tau_n c_n$ on w that has no trailing ε -transitions, i.e.

1. $n = 0$ if $w = \varepsilon$ (which implies that c_0 is accepting), and
2. $\ell(\tau_0 \cdots \tau_{n-1})$ is a strict prefix of w , if $w \neq \varepsilon$,

and $\tau_{n'} = r(\tau_0 \cdots \tau_{n'-1}, w(|\ell(\tau_0 \cdots \tau_{n'-1})|))$ for all $0 \leq n' < n$. If w is nonempty, $w(|\ell(\tau_0 \cdots \tau_{n'-1})|)$ is defined for all $0 \leq n' < n$ by the second requirement. Note that ρ is unique if it exists.

Note that the prefix processed so far can be recovered from r 's input, i.e. it is $\ell(\rho)$. However, the converse is not true due to the existence of ε -transitions. This is the reason that the run prefix and not the input prefix is the argument for the resolver. We denote the class of languages recognised by GFG-PDA by GFG-CFL.

Intuitively, every DPDA *should* be good-for-games, as there is no nondeterminism to resolve during a run. However, in order to reach a final state, a run of a DPDA on some input w may traverse *trailing* ε -transitions after the last letter of w is processed. On the other hand, the run of a GFG-PDA on w consistent with any resolver has to end with the transition processing the last letter of w . Hence, not every DPDA recognises the same language when viewed as a GFG-PDA. Nevertheless, we show, using standard pushdown automata constructions, that every DPDA can be turned into an equivalent GFG-PDA. As every GFG-PDA is a PDA by definition, we obtain a hierarchy of languages.

Lemma 1. $DCFL \subseteq GFG-CFL \subseteq CFL$.

Instead of requiring that GFG-PDA end their run with the last letter processed, one could add an end-of-word marker that allows traversing trailing ε -transitions after the last letter has been processed. In Appendix A.1, we show that this alternative definition does not increase expressiveness, which explains our (arguably simpler) definition.

Finally, let us remark that GFGness of PDA and context-free languages is undecidable. These problems were shown to be undecidable for ω -GFG-PDA and ω -GFG-CFL by reductions from the inclusion and universality problem for PDA on finite words [25]. Similar proofs also show that these problems are undecidable over PDA on finite words.

Theorem 1. *The following problems are undecidable:*

1. *Given a PDA \mathcal{P} , is \mathcal{P} a GFG-PDA?*
2. *Given a PDA \mathcal{P} , is $L(\mathcal{P}) \in GFG-CFL$?*

2.3 Games and Universality

One of the motivations for GFG automata is that solving games with winning conditions given by a GFG automaton is easier than for nondeterministic automata. This makes them appealing for applications such as the synthesis of reactive systems, which can be modelled as a game between an antagonistic environment and the system. Solving games is undecidable for PDA in general [12], both over finite and infinite words, while for ω -GFG-PDA, it is EXPTIME-complete [25]. As a corollary, universality is also decidable for ω -GFG-PDA, while it is undecidable for PDA, both over finite and infinite words [19].

Here, we consider Gale-Stewart games [13], abstract games induced by a language in which two players alternately pick letters, thereby constructing an infinite word. One player aims to construct a word that is in the language while the other aims to construct one that is not in the language. Note that these games are different, but related, to games played on configuration graphs of pushdown automata [40].

Formally, given a language $L \subseteq (\Sigma_1 \times \Sigma_2)^*$ of sequences of letter pairs, the game $G(L)$ is played between Player 1 and Player 2 in rounds $i = 0, 1, \dots$ as follows: At each round i , Player 1 plays a letter $a_i \in \Sigma_1$ and Player 2 answers with a letter $b_i \in \Sigma_2$. A play of $G(L)$ is an infinite word $\binom{a_0}{b_0} \binom{a_1}{b_1} \cdots$ and Player 2 wins such a play if and only if each of its prefixes is in the language L . A strategy for Player 2 is a mapping from Σ_1^+ to Σ_2 that gives for each prefix played by Player 1 the next letter to play. A play agrees with a strategy σ if for each i , $b_i = \sigma(a_0 a_1 \dots a_i)$. Player 2 wins $G(L)$ if she has a strategy that only agrees with plays that are winning for Player 2. Observe that Player 2 loses whenever the projection of L onto its first component is not universal. Finally, universality reduces to solving these games: \mathcal{P} is universal if and only if Player 2 wins $G(L)$ for $L = \{ \binom{w}{\#|w|} \mid w \in L(\mathcal{P}) \}$.

We now argue that solving games for GFG-PDA easily reduces to the case of ω -GFG-PDA, which are just GFG-PDA over infinite words, where acceptance is not determined by final state, since runs are infinite, but rather by the states or transitions visited infinitely often. Here, we only need safety ω -GFG-PDA, in which every infinite run is accepting (i.e. rejection is implemented via missing transitions). The infinite Gale-Stewart game over a language L of infinite words, also denoted by $G(L)$, is as above, except that victory is determined by whether the infinite word built along the play is in L .

Lemma 2. *Given a GFG-PDA \mathcal{P} , there is a safety ω -GFG-PDA \mathcal{P}' no larger than \mathcal{P} such that Player 2 wins $G(L(\mathcal{P}))$ if and only if she wins $G(L(\mathcal{P}'))$.*

Our main results of this section are now direct consequences, as argued above.

Corollary 1. *Given a GFG-PDA \mathcal{P} , deciding whether $L(\mathcal{P}) = \Sigma^*$ and whether Player 2 wins $G(L(\mathcal{P}))$ are both in EXPTIME.*

2.4 Closure properties

Like ω -GFG-PDA, GFG-PDA have poor closure properties.

Theorem 2. *GFG-PDA are not closed under union, intersection, complementation, set difference and homomorphism.*

The proofs are similar to those used for ω -GFG-PDA and relegated to the appendix. There, we also study the closure properties under these operations with regular languages: If L is in GFG-CFL and R is regular, then $L \cup R$, $L \cap R$ and $L \setminus R$ are also in GFG-CFL, but $R \setminus L$ is not necessarily in GFG-CFL.

3 Expressiveness

Here we show that GFG-PDA are more expressive than DPDA but less expressive than PDA.

Theorem 3. *$DCFL \subsetneq GFG-CFL \subsetneq CFL$.*

To show that GFG-PDA are more expressive than deterministic ones, we consider the language $B_2 = \{a^i \$ a^j \$ b^k \$ \mid k \leq \max(i, j)\}$. It is recognised by the PDA \mathcal{P}_{B_2} depicted in Figure 2, hence $B_2 \in CFL$. The first two states q_1 and q_2 deterministically push the input onto the stack, until the occurrence of the second $\$$. When the second $\$$ is processed, there is a nondeterministic choice to move to p_1 or p_2 and erase along ε -transitions 1 or 0 blocks from the stack, so that the 1st or 2nd block of a 's respectively remains at the top of the stack. Then, the automaton compares the length of the b -block in the input with the length of the a -block at the top of the stack and accepts if the b -block is shorter, i.e. the third $\$$ is processed before the whole a -block is popped off the stack. If the input has not the form $a^i \$ a^j \$ b^k \$$, then it is rejected.




Fig. 2. A PDA \mathcal{P}_{B_2} recognising B_2 . Grey states are final, and X is an arbitrary stack symbol.

We show that $B_2 \in GFG-CFL$ by proving that \mathcal{P}_{B_2} is good-for-games: the nondeterministic choice between moving to p_1 or to p_2 can be made only based on the prefix $a^i \$ a^j$ processed so far. This is straightforward, as a resolver only needs to know which of i and j is larger, which can be determined from the run prefix constructed thus far. Then, in order to show that B_2 is not in $DCFL$, we prove that its complement B_2^c is not a context-free language. Since $DCFL$ is closed under complementation, this implies the desired result.

Finally, to show that PDA are more expressive than GFG-PDA, we consider the language $L = \{a^n b^n \mid n \geq 0\} \cup \{a^n b^{2n} \mid n \geq 0\}$. We note that $L \in CFL$ while we show below $L \notin GFG-CFL$. All proofs can be found in Appendix A.3.

Unambiguous context-free languages, i.e. those generated by grammars for which every word in the language has a unique leftmost derivation, are another class sitting between $DCFL$ and CFL . Thus, it is natural to ask how unambiguity and GFGness are related: To conclude this section, we show that both notions are independent.

Theorem 4. *There is an unambiguous context-free language that is not in GFG-CFL and a language in GFG-CFL that is inherently ambiguous.*

An unambiguous grammar for the language $\{a^n b^n \mid n \geq 0\} \cup \{a^n b^{2n} \mid n \geq 0\} \notin \text{GFG-CFL}$ is easy to construct and we show in Appendix A.3 that the language $B = \{a^i b^j c^k \mid i, j, k \geq 1, k \leq \max(i, j)\}$ is inherently ambiguous. Its inclusion in GFG-CFL is easily established using a similar argument as for the language $B_2 = \{a^i \$ a^j \$ b^k \$ \mid k \leq \max(i, j)\}$ above. The dollars add clarity to the GFG-PDA but are cumbersome in the proof of inherent ambiguity.

4 Succinctness

We show that GFG-PDA are not only more expressive than DPDA, but also more succinct. Similarly, we show that PDA are more succinct than GFG-PDA.

Theorem 5. *GFG-PDA can be exponentially more succinct than DPDA, and PDA can be double-exponentially more succinct than GFG-PDA.*

We first show that GFG-PDA can be exponentially more succinct than DPDA. To this end, we construct a family $(C_n)_{n \in \mathbb{N}}$ of languages such that C_n is recognised by a GFG-DPDA of size $O(n)$, yet every DPDA recognising C_n has at least exponential size in n .

Let $c_n \in (\{0, 1\}^n)^*$ be the word describing an n -bit binary counter counting from 0 to $2^n - 1$. For example, $c_2 = \$00\$01\$10\11 . We consider the family of languages $C_n = \{w \in \{0, 1, \$, \#\}^* \mid w \neq c_n \#\} \subseteq \{0, 1, \$, \#\}^*$ of bad counters.

We show in Appendix A.3 that the language C_n is recognised by a GFG-PDA of size $O(n)$ and that every DPDA \mathcal{D} recognising C_n has exponential size in n . Observe that this result implies that even GFG-PDA that are equivalent to DPDA are not determinisable by pruning. In contrast, for NFA GFGness implies determinisability by pruning [6].

We conclude this section by showing that PDA can be double-exponentially more succinct than GFG-PDA. We show that there exists a family $(L_n)_{n > 0}$ of languages such that L_n is recognised by a PDA of size $O(\log n)$ while every GFG-PDA recognising this language has at least exponential size in n .

Formally, we set $L_n = (0 + 1)^* 1 (0 + 1)^{n-1}$, that is, the n^{th} bit from the end is a 1. We count starting from 1, so that the last bit is the 1st bit from the end. Note that this is the standard example for showing that NFA can be exponentially more succinct than DFA, and has been used for many other succinctness results ever since.

In Appendix A.3, we first show that L_n is recognised by a PDA of size $O(\log n)$. To conclude, we prove that every GFG-PDA recognising L_n has at least exponential size in n .

5 Good-for-games Visibly Pushdown Automata

One downside of GFG-PDA is that, like ω -GFG-PDA, they have poor closure properties and checking GFGness is undecidable. We therefore consider a well-behaved class of GFG-PDA, namely GFG visibly pushdown automata, GFG-VPA for short, that is closed under union, intersection, and complementation.

Let Σ_c, Σ_r and Σ_{int} be three disjoint sets of *call* symbols, *return* symbols and *internal* symbols respectively. Let $\Sigma = \Sigma_c \cup \Sigma_r \cup \Sigma_{\text{int}}$. A *visibly pushdown automaton* [2] (VPA) $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ is a restricted PDA that pushes onto the stack only when it reads a call symbol, it pops the stack only when a return symbol is read, and does not use the stack when there is an internal symbol. Formally,

- 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 $Y \in \Gamma$ 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 $a \in \Sigma_{\text{int}}$ 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 $w \in (\Sigma_c \cup \Sigma_r \cup \Sigma_s)^*$ only depends on w .

We denote by GFG-VPA the VPA that are good-for-games. Every VPA (and hence every GFG-VPA) can be determinised, i.e. all three classes of automata recognise the same class of languages, denoted by VPL, which is a strict subset of DCFL [2]. However, VPA can be exponentially more succinct than deterministic VPA (DVPA) [2].

Theorem 6. *GFG-VPA can be exponentially more succinct than DVPA and VPA can be exponentially more succinct than GFG-VPA.*

GFGness of VPA is decidable using the *one-token game*, introduced by Bagnol and Kuperberg [3]. It modifies the game-based characterisation of GFGness of ω -regular automata by Henzinger and Piterman [16]. While the one-token game does not characterise the GFGness of Büchi automata, here we show that it suffices for VPA. The matching lower bound follows from a reduction from the inclusion problem for VPA, which is EXPTIME-hard [2], to GFGness (see [25] for details of the reduction in the context of ω -GFG-PDA).

Theorem 7. *The following problem is EXPTIME-complete: Given a VPA \mathcal{P} , is \mathcal{P} GFG?*

Finally, we relate the GFGness problem to the *good-enough synthesis* problem [1], also known as the *uniformization* problem [9], which is similar to the Church synthesis problem, except that the system is only required to satisfy the specification on inputs in the projection of the specification on the first component.

Definition 1 (GE-synthesis). *Given a language $L \subseteq (\Sigma_1 \times \Sigma_2)^*$, is there a function $f : \Sigma_1^* \rightarrow \Sigma_2$ such that for each $w \in \{w \mid \exists w' \in \Sigma_2^*. (w, w') \in L\}$ the word $(w, f(w(0))f(w(0)w(1)) \dots)$ is in L ?*

Corollary 2. *The GE-synthesis problem for input given as GFG-VPA, and in particular for DVPA, is EXPTIME-complete*

In contrast, for LTL specifications, the GE-synthesis problem is 2EXPTIME-complete [1].

6 Resolvers

The definition of a resolver does not put any restrictions on its complexity. In this section we study the complexity of the resolvers that GFG-PDA need. We consider two somewhat orthogonal notions of complexity: memory and machinery. On one hand, we show that resolvers can always be chosen to be *positional*, that is, dependent on the current state and stack configuration only. Note that this is not the case for ω -regular automata², let alone ω -GFG-PDA. On the other hand, we show that they are not always implementable by pushdown transducers.

More formally, a resolver r is positional, if for any two sequences ρ and ρ' of transitions inducing runs ending in the same configuration, $r(\rho, a) = r(\rho', a)$ for all $a \in \Sigma$.

Lemma 3. *Every GFG-PDA has a positional resolver.*

Contrary to the case of finite and ω -regular automata, since GFG-PDA have an infinite configuration space, the existence of positional resolvers does not imply determinisability. On the other hand, if a GFG-PDA has a resolver which only depends on the mode of the current configuration, then it is *determinisable by pruning*, as transitions that are not used by the resolver can be removed to obtain a deterministic automaton. However, not all GFG-PDA are determinisable by pruning, e.g. the GFG-PDA for the languages C_n used to prove Theorem 5.

We now turn to how powerful resolvers for GFG-PDA need to be. First, we introduce transducers as a way to implement a resolver. A transducer is an automaton with outputs instead of acceptance, i.e., it computes a function from input sequences to outputs. A pushdown resolver is a pushdown transducer that implements a resolver.

² A positional resolver for ω -regular automata implies determinisability by pruning, and we know that this is not always possible [6]

Note that a resolver has to pick enabled transitions in order to induce accepting runs for all inputs in the language. To do so, it needs access to the mode of the last configuration. However, to keep track of this information on its own, the pushdown resolver would need to simulate the stack of the GFG-PDA it controls. This severely limits the ability of the pushdown resolver to implement computations on its own stack. Thus, we give a pushdown resolver access to the current mode of the GFG-PDA via its output function, thereby freeing its own stack to implement further functionalities.

Formally, a pushdown transducer (PDT for short) $\mathcal{T} = (\mathcal{D}, \lambda)$ consists of a DPDA \mathcal{D} augmented with an output function $\lambda : Q^{\mathcal{D}} \rightarrow \Theta$ mapping the states $Q^{\mathcal{D}}$ of \mathcal{D} to an output alphabet Θ . The input alphabet of \mathcal{T} is the input alphabet of \mathcal{D} .

Then, given a PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$, a pushdown resolver for \mathcal{P} consists of a pushdown transducer $\mathcal{T} = (\mathcal{D}, \lambda)$ with input alphabet Δ and output alphabet $Q \times \Gamma_{\perp} \times \Sigma \rightarrow \Delta$ such that the function $r_{\mathcal{T}}$, defined as follows, is a resolver for \mathcal{P} : $r_{\mathcal{T}}(\tau_0 \dots \tau_k, a) = \lambda(q_{\mathcal{T}})(q_{\mathcal{P}}, X, a)$ where

- $q_{\mathcal{T}}$ is the state of the last configuration of the longest run of \mathcal{D} on $\tau_0 \dots \tau_k$ (recall that while \mathcal{D} is deterministic, it may have several runs on an input which differ on trailing ε -transitions);
- $(q_{\mathcal{P}}, X)$ is the mode of the last configuration of the run of \mathcal{P} induced by $\tau_0 \dots \tau_k$.

In other words, a transducer implements a resolver by processing the run so far, and then uses the output of the state reached and the state and top stack symbol of the GFG-PDA to determine the next transition in the GFG-PDA.




Fig. 3. The PDA \mathcal{P}_{B_3} for B_3 . Grey states are final, and X is an arbitrary stack symbol.

We now give an example of a GFG-PDA which does not have a pushdown resolver. The language in question is the language $B_3 = \{a^i \$ a^j \$ a^k \$ b^l \$ \mid l \leq \max(i, j, k)\}$. Compare this to the language B_2 in Section 3 which *does* have a pushdown resolver. Let \mathcal{P}_{B_3} be the automaton in Figure 3, which works analogously to the automaton for B_2 in Figure 2.

This automaton recognises B_3 : for a run to end in the final state, the stack, and therefore the input, must have had an a -block longer than or equal to the final b -block; conversely, if the b -block is shorter than or equal to some a -block, the automaton can nondeterministically pop the blocks on top of the longest a -block off the stack before processing the b -block. Furthermore, this automaton is GFG: the nondeterminism can be resolved by removing from the stack all blocks until the *longest* a -block is at the top of the stack, and this choice can be made once the third $\$$ is processed.

We now argue that this GFG-PDA needs more than a pushdown resolver.

Lemma 4. *The GFG-PDA \mathcal{P}_{B_3} has no pushdown resolver.*

Another restricted class of resolvers are finite-state resolvers, which can be seen as pushdown resolvers that do not use their stack. Similarly to the case of ω -GFG-PDA [26], the product of a GFG-PDA and a finite-state resolver yields a DPDA for the same language.

Remark 2. Every GFG-PDA with a finite-state resolver is determinisable.

Note that the converse does not hold. For example, consider the regular, and therefore deterministic context-free, language L_{10} of words $w\#$ with $w \in \{a, b\}^*$ with infix a^{10} . A GFG-PDA \mathcal{P}_{10} recognising

L_{10} can be constructed as follows: \mathcal{P}_{10} pushes its input onto its stack until processing the first $\#$. Then, it uses ε -transitions to empty the stack again. While doing so, it can nondeterministically guess whether the next 10 letters removed from the stack are all a 's. If yes, it accepts; in all other cases (in particular if the input word does not end with the first $\#$ or the infix a_{10} is not encountered on the stack) it rejects. This automaton is good-for-games, as a resolver has access to the whole prefix before the first $\#$ when searching for a^{10} while emptying the stack. This is sufficient to resolve the nondeterminism. On the other hand, there is no finite-state resolver for \mathcal{P}_{10} , as resolving the nondeterminism, intuitively, requires to keep track of the whole prefix before the first $\#$ (recall that a finite-state resolver only has access to the topmost stack symbol).

In Appendix A.2 we consider another model of pushdown resolver, namely one that does not only have access to the mode of the GFG-PDA, but can check the full stack for regular properties. We show that this change does not increase the class of good-for-games context-free languages that are recognised by a GFG-PDA with a pushdown resolver.

Finally, for GFG-VPA, the situation is again much better. The classical game-based characterisation of GFGness of ω -regular automata by Henzinger and Piterman [16] can be lifted to VPA and is, crucially, decidable.

Theorem 8. *Every GFG-VPA has a (visibly) pushdown resolver.*

7 Conclusion

We have continued the study of good-for-games pushdown automata, focusing on expressiveness and succinctness. In particular, we have shown that GFG-PDA are not only more expressive than DPDA (as had already been shown for the case of infinite words), but also more succinct than DPDA: We have introduced the first techniques for using GFG nondeterminism to succinctly represent languages that do not depend on the coBüchi condition. Similarly, for the case of VPA, for which deterministic and nondeterministic automata are equally expressive, we proved a (tight) exponential gap in succinctness.

Solving games and universality are decidable for GFG-PDA, but GFGness is undecidable and GFG-PDA have limited closure properties. On the other hand, GFGness for VPA is decidable and they inherit the closure properties of VPA, e.g. union, intersection and complementation, making GFG-VPA an exciting class of pushdown automata. Finally, we have studied the complexity of resolvers for GFG-PDA, showing that positional ones always suffice, but that they are not always implementable by pushdown transducers. Again, GFG-VPA are better-behaved, as they always have a resolver implementable by a VPA.

Let us conclude by mentioning some open problems raised by our work.

- It is known that the succinctness gap between PDA and DPDA is noncomputable, i.e. there is no computable function f such that any PDA of size n that has some equivalent DPDA also has an equivalent DPDA of size $f(n)$. Due to our hierarchy results, at least one of the succinctness gaps between PDA and GFG-PDA and between GFG-PDA and DPDA has to be uncomputable, possibly both.
- We have shown that some GFG-PDA do not have pushdown resolvers. It is even open whether every GFG-PDA has a computable resolver.
- On the level of languages, it is open whether every language in GFG-CFL has a GFG-PDA recognising it with a resolver implementable by a pushdown transducer.
- We have shown that GFGness is undecidable, both for PDA and for context-free languages. Is it decidable whether a given GFG-PDA has an equivalent DPDA?
- Equivalence of DPDA is famously decidable [37] while it is undecidable for PDA [19]. Is equivalence of GFG-PDA decidable?
- Does every GFG-PDA that is equivalent to a DPDA have a finite-state resolver with regular stack access (see Appendix A.2 for definitions)?
- There is a plethora of fragments of context-free languages one can compare GFG-CFL to, let us just mention a few interesting ones: Height-deterministic context-free languages [31], context-free languages with bounded nondeterminism [17] and preorder typeable visibly pushdown languages [21].

Acknowledgements



This work is part of projects that receive funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreements No 754411 and No 892704.

References

1. Almagor, S., Kupferman, O.: Good-enough synthesis. In: Lahiri, S.K., Wang, C. (eds.) CAV 2020, Part II. LNCS, vol. 12225, pp. 541–563. Springer (2020), https://doi.org/10.1007/978-3-030-53291-8_28
2. Alur, R., Madhusudan, P.: Visibly pushdown languages. In: Babai, L. (ed.) STOC 2004. pp. 202–211. ACM (2004)
3. Bagnol, M., Kuperberg, D.: Büchi Good-for-Games Automata Are Efficiently Recognizable. In: Ganguly, S., Pandya, P. (eds.) FSTTCS 2018. LIPIcs, vol. 122, pp. 16:1–16:14. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2018), <http://drops.dagstuhl.de/opus/volltexte/2018/9915>
4. Björklund, H., Martens, W.: The tractability frontier for NFA minimization. *Journal of Computer and System Sciences* 78(1), 198–210 (2012)
5. Björklund, H., Schuster, M., Schwentick, T., Kulbatzki, J.: On optimum left-to-right strategies for active context-free games. In: Tan, W., Guerrini, G., Catania, B., Gounaris, A. (eds.) ICDT 2013. pp. 105–116. ACM (2013), <https://doi.org/10.1145/2448496.2448510>
6. Boker, U., Kupferman, O., Skrzypczak, M.: How deterministic are good-for-games automata? arXiv:1710.04115 (2017)
7. Büchi, J.R.: Regular canonical systems. *Archiv für mathematische Logik und Grundlagenforschung* 6(3), 91–111 (Apr 1964)
8. Carayol, A., Hague, M.: Saturation algorithms for model-checking pushdown systems. In: Ésik, Z., Fülöp, Z. (eds.) AFL 2014. EPTCS, vol. 151, pp. 1–24 (2014)
9. Carayol, A., Löding, C.: Uniformization in automata theory. In: LMPS 2015 (2015), <https://hal.archives-ouvertes.fr/hal-01806575>
10. Charikar, M., Lehman, E., Liu, D., Panigrahy, R., Prabhakaran, M., Sahai, A., Shelat, A.: The smallest grammar problem. *IEEE Trans. Inf. Theory* 51(7), 2554–2576 (2005), <https://doi.org/10.1109/TIT.2005.850116>
11. Colcombet, T.: The theory of stabilisation monoids and regular cost functions. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikolettseas, S.E., Thomas, W. (eds.) ICALP 2009, (Part II). LNCS, vol. 5556, pp. 139–150. Springer (2009)
12. Finkel, O.: Topological properties of omega context-free languages. *Theor. Comput. Sci.* 262(1), 669–697 (2001)
13. Gale, D., Stewart, F.M.: Infinite games with perfect information. In: Kuhn, H.W., Tucker, A.W. (eds.) *Contributions to the Theory of Games (AM-28)*, Volume II, chap. 13, pp. 245–266. Princeton University Press (1953)
14. Goldstine, J., Leung, H., Wotschke, D.: Measuring nondeterminism in pushdown automata. *Journal of Computer and System Sciences* 71(4), 440–466 (2005)
15. Hartmanis, J.: On the succinctness of different representations of languages. *SIAM J. Comput.* 9(1), 114–120 (1980)
16. Henzinger, T.A., Piterman, N.: Solving games without determinization. In: Ésik, Z. (ed.) CSL 2006. LNCS, vol. 4207, pp. 395–410. Springer (2006)
17. Herzog, C.: Pushdown automata with bounded nondeterminism and bounded ambiguity. *Theor. Comput. Sci.* 181(1), 141–157 (1997)
18. Holzer, M., Kutrib, M.: Descriptive complexity of (un)ambiguous finite state machines and pushdown automata. In: Kucera, A., Potapov, I. (eds.) RP 2010. LNCS, vol. 6227, pp. 1–23. Springer (2010)
19. Hopcroft, J.E., Ullman, J.D.: *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley (1979)
20. Jiang, T., Ravikumar, B.: Minimal NFA problems are hard. *SIAM Journal on Computing* 22(6), 1117–1141 (1993)
21. Kumar, V., Madhusudan, P., Viswanathan, M.: Visibly pushdown automata for streaming XML. In: Williamson, C.L., Zurko, M.E., Patel-Schneider, P.F., Shenoy, P.J. (eds.) WWW 2007. pp. 1053–1062. ACM (2007), <https://doi.org/10.1145/1242572.1242714>
22. Kuperberg, D., Majumdar, A.: Computing the width of non-deterministic automata. *Log. Methods Comput. Sci.* 15(4) (2019)
23. Kuperberg, D., Skrzypczak, M.: On determinisation of good-for-games automata. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) ICALP 2015 (Part II). LNCS, vol. 9135, pp. 299–310. Springer (2015)

24. Kupferman, O., Safra, S., Vardi, M.Y.: Relating word and tree automata. *Annals of Pure and Applied Logic* 138(1-3), 126–146 (2006)
25. Lehtinen, K., Zimmermann, M.: Good-for-games ω -pushdown automata. In: Hermanns, H., Zhang, L., Kobayashi, N., Miller, D. (eds.) *LICS 2020*. pp. 689–702. ACM (2020)
26. Lehtinen, K., Zimmermann, M.: Good-for-games ω -pushdown automata. *arXiv 2001.04392* (2020)
27. Löding, C., Madhusudan, P., Serre, O.: Visibly pushdown games. In: Lodaya, K., Mahajan, M. (eds.) *FSTTCS 2004*. LNCS, vol. 3328, pp. 408–420. Springer (2004)
28. Martens, W., Neven, F., Schwentick, T., Bex, G.J.: Expressiveness and complexity of XML schema. *TODS* 31(3), 770–813 (2006)
29. Maurer, H.A.: A direct proof of the inherent ambiguity of a simple context-free language. *J. ACM* 16(2), 256–260 (1969), <https://doi.org/10.1145/321510.321517>
30. Muscholl, A., Schwentick, T., Segoufin, L.: Active context-free games. *Theory of Computing Systems* 39(1), 237–276 (2006)
31. Nowotka, D., Srba, J.: Height-deterministic pushdown automata. In: Kucera, L., Kucera, A. (eds.) *MFCSS 2007*. LNCS, vol. 4708, pp. 125–134. Springer (2007), https://doi.org/10.1007/978-3-540-74456-6_13
32. Okhotin, A., Salomaa, K.: Descriptive complexity of unambiguous input-driven pushdown automata. *Theor. Comput. Sci.* 566, 1–11 (2015)
33. Radi, B.A., Kupferman, O.: Minimizing GFG Transition-Based Automata. In: Baier, C., Chatzigiannakis, I., Flocchini, P., Leonardi, S. (eds.) *ICALP 2019*. LIPIcs, vol. 132, pp. 100:1–100:16. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2019)
34. Salomaa, A., Soittola, M.: *Automata-Theoretic Aspects of Formal Power Series*. Texts and Monographs in Computer Science, Springer (1978), <https://doi.org/10.1007/978-1-4612-6264-0>
35. Salomaa, K., Yu, S.: Limited nondeterminism for pushdown automata. *Bull. EATCS* 50, 186–193 (1993)
36. Schewe, S.: Minimising Good-For-Games Automata Is NP-Complete. In: Saxena, N., Simon, S. (eds.) *FSTTCS 2020*. LIPIcs, vol. 182, pp. 56:1–56:13. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2020)
37. Sénizergues, G.: $L(A)=L(B)$? decidability results from complete formal systems. *Theor. Comput. Sci.* 251(1-2), 1–166 (2001)
38. Stearns, R.E., Hunt III, H.B.: On the equivalence and containment problems for unambiguous regular expressions, regular grammars and finite automata. *SIAM Journal on Computing* 14(3), 598–611 (1985)
39. Valiant, L.G.: A note on the succinctness of descriptions of deterministic languages. *Inf. Control.* 32(2), 139–145 (1976)
40. Walukiewicz, I.: Pushdown processes: Games and model-checking. *Inf. Comput.* 164(2), 234–263 (2001)

A Appendix

A.1 Resolvers with End-of-word Markers

As mentioned in the main part, GFG-PDA are by definition required to end their run with the last letter of the input word. Instead, one could also consider a model where they are allowed to take some trailing ε -transitions after the last input letter has been processed. As a resolver has access to the next input letter, which is undefined in this case, we need resolvers with end-of-word markers to signal the resolver that the last letter has been processed. In the following, we show that GFG-PDA with end-of-word resolvers are as expressive as standard GFG-PDA, albeit exponentially more succinct.

Fix some distinguished end-of-word-marker $\#$, which takes the role of the next input letter to be processed, if there is none after the last letter of the input word is processed. Let $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ be a PDA with $\# \notin \Sigma$. An EoW-resolver for \mathcal{P} is a function $r: \Delta^* \times (\Sigma \cup \{\#\}) \rightarrow \Delta$ such that for every $w \in L(\mathcal{P})$, there is an accepting run $c_0 \tau_0 \cdots \tau_n c_n$ on w such that $\tau_{n'} = r(\tau_0 \cdots \tau_{n'-1}, w\#(|\ell(\tau_0 \cdots \tau_{n'-1})|))$ for all $0 \leq n' < n$. Note that the second argument given to the resolver is a letter of $w\#$, which is equal to $\#$ if the run prefix induced by $\tau_0 \cdots \tau_{n'-1}$ has already processed the full input w . Again, the run is unique if it exists and, but may have trailing ε -transitions.

Lemma 5. *GFG-PDA with EoW-resolvers are as expressive as GFG-PDA.*

Proof. A (standard) resolver can be turned into an EoW-resolver that ignores the EoW-marker. Hence, every GFG-PDA is a PDA with EoW-resolver recognizing the same language. So, it only remains to consider the other inclusion.

To this end, let $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ be a PDA with EoW-resolver. The language

$$C_{acc} = \{\gamma q \mid q \in F \text{ and } \gamma \in \perp \Gamma^*\} \subseteq \perp \Gamma^* Q$$

encoding final configurations of \mathcal{P} is regular. Hence, the language

$$C = \{\gamma q \in \perp \Gamma^* Q \mid \text{there is a run infix } (q, \gamma) \tau_0 \cdots \tau_{n-1} c_n \\ \text{with } \ell(\tau_0 \cdots \tau_{n-1}) = \varepsilon \text{ and } c_n \in C_{acc}\}$$

can be shown to be regular as well by applying saturation techniques [7]³ to the restriction of \mathcal{P} to ε -transitions. If \mathcal{P} reaches a configuration $c \in C$ after processing an input w , then $w \in L$, even if c 's state is not final.

Let $\mathcal{A} = (Q_{\mathcal{A}}, \Gamma_{\perp} \cup Q, q_I^{\mathcal{A}}, \delta_{\mathcal{A}}, F_{\mathcal{A}})$ be a DFA recognizing C . We extend the stack alphabet of \mathcal{P} to $\Gamma \times Q_{\mathcal{A}} \times (Q_{\mathcal{A}} \cup \{u\})$, where u is a fresh symbol. Then, we extend the transition relation such that it keeps track of the unique run of \mathcal{A} on the stack content: If \mathcal{P} reaches a stack content $\perp(X_1, q_1, q'_1)(X_2, q_2, q'_2) \cdots (X_s, q_s, q'_s)$, then we have

$$q_j = \delta_{\mathcal{A}}^*(q_I^{\mathcal{A}}, \perp X_1 \cdots X_j)$$

for every $1 \leq j \leq s$ as well as $q'_j = q'_{j-1}$ for every $2 \leq j \leq s$ and $q'_1 = u$. Here, $\delta_{\mathcal{A}}^*$ is the standard extension of $\delta_{\mathcal{A}}$ to words. The adapted PDA is still good-for-games, as no new nondeterminism has been introduced, and keeps track of the state of \mathcal{A} reached by processing the stack content as well as the shifted sequence of states of \mathcal{A} , which is useful when popping the top stack symbol: If the topmost stack symbol (X, q, q') is popped of the stack then q' is the state of \mathcal{A} reached when processing the remaining stack.

Now, we double the state space of \mathcal{P} , making one copy final, and adapt the transition relation again so that a final state is reached whenever \mathcal{P} would reach a configuration in C . Whether a configuration in C is reached can be determined from the current state of \mathcal{P} being simulated, as well as the top stack symbol containing information on the run of \mathcal{A} on the current stack content. The resulting PDA recognizes $L(\mathcal{P})$ and has on every word $w \in L(\mathcal{P})$ an accepting run without trailing ε -transitions. Furthermore, an EoW-resolver for \mathcal{P} can be turned into a (standard) resolver for \mathcal{P}' , as the tracking of stack contents and the doubling of the state space does not introduce nondeterminism. \square

³ Also, see the survey by Carayol and Hague [8] for more details.

As \mathcal{A} has at most exponential size, \mathcal{P} is also exponential (both in the size of \mathcal{P}). This exponential blowup incurred by removing the end-of-word marker is in general unavoidable. In Lemma 20, we show that the language L_n of bitstrings whose n^{th} bit from the end is a 1 requires exponentially-sized GFG-PDA. On the other hand, it is straightforward to devise a polynomially-sized PDA with EoW-marker recognizing L_n : the underlying PDA stores the input word on the stack, guesses nondeterministically that the word has ended, uses n (trailing) ε -transitions to pop of the last $n - 1$ letters stored on the stack, and then checks that the topmost stack symbol is an 1. With an EoW-resolver, the end of the input does not have to be guessed, but is marked by the EoW-marker.

A.2 Pushdown Resolvers with Regular Stack Access

Recall that pushdown transducers implementing a resolver have access to the mode of the GFG-PDA whose nondeterminism it resolves. Here, we consider a more general model where the transducer can use information about the whole stack when determining the next transition. More precisely, we consider a regular abstraction of the possible stack contents by fixing a DFA running over the stack and allowing the transducer to base its decision on the state reached by the DFA as well.

Then, given a PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$, a pushdown resolver with regular stack access $\mathcal{T} = (\mathcal{D}, \mathcal{A}, \lambda)$ consists a DPDA \mathcal{P} with input alphabet Δ , a DFA \mathcal{A} over Γ_{\perp} with state set $Q^{\mathcal{A}}$, and an output function λ with output alphabet $Q \times Q^{\mathcal{A}} \times \Sigma \rightarrow \Delta$ such that the function $r_{\mathcal{T}}$ defined as follows, is a resolver for \mathcal{P} :

$$r_{\mathcal{T}}(\tau_0 \dots \tau_k, a) = \lambda(q_{\mathcal{T}})(q_{\mathcal{P}}, q_{\mathcal{A}}, a)$$

where

- $q_{\mathcal{T}}$ is the state of the last configuration of the longest run of \mathcal{D} on $\tau_0 \dots \tau_k$ (recall that while \mathcal{D} is deterministic, it may have several runs on an input which differ on trailing ε -transitions).
- Let c be the last configuration of the run of \mathcal{P} induced by $\tau_0 \dots \tau_k$. Then, $q_{\mathcal{P}}$ is the state of c and $q_{\mathcal{A}}$ is the state of \mathcal{A} reached when processing the stack content of c .

Every pushdown resolver with only access to the current mode is a special case of a pushdown resolver with regular stack access. On the other hand, having regular access to the stack is strictly stronger than having just access to the mode. However, by adapting the underlying GFG-PDA, one can show that the *languages* recognised by GFG-PDA with pushdown resolvers do not increase when allowing regular stack access.

Lemma 6. *Every GFG-PDA with a pushdown resolver with regular stack access can be turned into an equivalent GFG-PDA with a pushdown resolver.*

Proof. Let $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ be a GFG-PDA and let $(\mathcal{D}, \mathcal{A}, \lambda)$ be a pushdown resolver with stack access for \mathcal{P} . We keep track of the state \mathcal{A} reaches on the current stack as in the proof of Lemma 5: If a stack content $\perp(X_1, q_1) \dots (X_s, q_s)$ is reached, then q_j is the unique state of \mathcal{P} reached when processing $\perp X_1 \dots X_j$. Now, it is straightforward to turn $(\mathcal{D}, \mathcal{A}, \lambda)$ into a pushdown resolver for \mathcal{P} that has only access to the top stack symbol. \square

A.3 Proofs Omitted due to Space Restrictions

Proof of Lemma 1 Recall that we need to prove $\text{DCFL} \subseteq \text{GFG-CFL} \subseteq \text{CFL}$.

Proof. We only consider the first inclusion, as the second one is trivial. So, let $L \in \text{DCFL}$, say it is recognised by the DPDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$. We say that a mode m of \mathcal{P} is a reading mode if it does not enable an ε -transition. Hence, due to determinism, m can only enable at most one a -transition for every $a \in \Sigma$.

Now, consider some nonempty word $w(0) \dots w(n) \in L(\mathcal{P})$ (we take care of the empty word later on), say with accepting run ρ (treated, for notational convenience, as a sequence of transitions). This run can be decomposed as

$$\rho = \rho_0 \tau_0 \rho_1 \tau_1 \rho_2 \dots \rho_n \tau_n \rho_{n+1}$$

where τ_i processes $w(i)$ and each ρ_i is a (possibly empty) sequence of ε -transitions. Each run prefix induced by some $\rho_0 \tau_0 \rho_1 \dots \rho_i$ ends in a configuration with reading mode.

Intuitively, we have to eliminate the trailing ε -transitions in ρ_{n+1} . To do so, we postpone the processing of the letter $w(i)$ to the end of ρ_{i+1} . Instead, we guess that the next input is $w(i)$ by turning the original $w(i)$ -transition τ_i into an ε -transition τ_i^d that stores $w(i)$ in the state space of the modified automaton. Then, ρ_{i+1} is simulated, a dummy transition τ_i^d processing the stored letter $w(i)$ is executed.

Hence, the resulting run of the modified automaton on w has the form

$$\rho_0 \tau_0' \rho_1 \tau_0^d \tau_1' \rho_2 \tau_1^d \cdots \rho_n \tau_{n-1}^d \tau_n' \rho_{n+1} \tau_n^d,$$

where each τ_i^d is now an ε -transition, each ρ_i is a (possibly empty) sequence of ε -transitions, and each τ_i^d is a dummy transition processing $w(i)$. Hence, the run ends with the transition processing the last letter $w(n)$ of the input.

The resulting PDA is good-for-games, as a resolver has access to the next letter to be processed, which is sufficient to resolve the nondeterminism introduced by the guessing of the next letter.

More formally, let PDA $\mathcal{P}' = (Q', \Sigma, \Gamma, q_I', \Delta', F')$ where

- $Q' = Q \cup (Q \times \Sigma)$,
- $q_I' = q_I$,
- $F' = F \cup I$ where $I = \{q_I\}$ if $\varepsilon \in L(\mathcal{P})$ and $I = \emptyset$, otherwise, and
- Δ' is the union of the following sets of transitions:
 - $\{\tau \in \Delta \mid \ell(\tau) = \varepsilon\}$, which is used to simulate the leading sequence of ε -transitions before the first letter is processed by \mathcal{P} , i.e. the transitions in ρ_0 above.
 - $\{(q, X, \varepsilon, (q', a), \gamma) \mid (q, X, a, q', \gamma) \in \Delta \text{ and } a \in \Sigma\}$, which are used to guess and store the next letter to be processed.
 - $\{((q, a), X, \varepsilon, (q', a), \gamma) \mid (q, X, \varepsilon, q', \gamma) \in \Delta\}$, which are used to simulate ε -transitions after a letter has been guessed, but not yet processed (i.e. transitions in some ρ_i with $i > 0$).
 - $\{((q, a), X, a, q, X) \mid (q, X) \text{ is a reading mode}\}$, the dummy transitions used to actually process the guessed and stored letter.

Now, formalising the intuition given above, one can show that \mathcal{P}' has a resolver witnessing that it recognises $L(\mathcal{P})$. In particular, the empty word is in $L(\mathcal{P}')$ if and only if it is in $L(\mathcal{P})$, as the run induced by the resolver on ε ends in the initial configuration, which is final if and only if $\varepsilon \in L(\mathcal{P})$. \square

Proof of Lemma 2 Recall that we need to prove the following statement: Given a GFG-PDA \mathcal{P} , there is a safety ω -GFG-PDA \mathcal{P}' no larger than \mathcal{P} such that Player 2 wins $G(L(\mathcal{P}))$ if and only if she wins $G(L(\mathcal{P}'))$.

Proof. Let \mathcal{P}' be the PDA obtained from \mathcal{P} by removing all transitions (q, X, a, q', γ) of \mathcal{P} with $a \in \Sigma$ and with non-final q' .

With a safety condition, in which every infinite run is accepting, \mathcal{P}' recognises exactly those infinite words whose prefixes are all accepted by \mathcal{P} . Hence, the games $G(L(\mathcal{P}))$ and $G(L(\mathcal{P}'))$ have the same winning player.

Note that the correctness of this construction crucially relies on our definition of GFG-PDA which requires a run on a finite word to end as soon as the last letter is processed. Hence, the word is accepted if and only if the state reached by processing this last letter is final.

Finally, since \mathcal{P} is GFG, so is \mathcal{P}' . Consider an infinite input in $L(\mathcal{P}')$. Then, every prefix w has an accepting run of \mathcal{P} induced by its resolver, which implies that the last transition of this run (which processes the last letter of w) is not one of those that are removed to obtain \mathcal{P}' . Now, an induction shows that the same resolver works for \mathcal{P}' as well, relying on the fact that if w and w' with $|w| < |w'|$ are two such prefixes, then the resolver-induced run of \mathcal{P} on w is a prefix of the resolver-induced run of \mathcal{P} on w' . \square

Proof of Theorem 2 We need to prove the non-closure of GFG-CFL under union, intersection, complementation, set difference and homomorphism. Some of the proofs refer to results proven later in the appendix.

Lemma 7. *GFG-CFL is not closed under union.*

Proof. Consider the languages $L_1 = \{a^n b^n \mid n \geq 0\}$ and $L_2 = \{a^n b^{2n} \mid n \geq 0\}$ respectively. There exist a DPDA recognising L_1 and a DPDA recognising L_2 . Hence by Lemma 1, there also exist a GFG-PDA recognising L_1 and a GFG-PDA recognising L_2 . However, by Lemma 14, we have that $L_1 \cup L_2$ cannot be recognised by a GFG-PDA. \square

Lemma 8. *GFG-CFL is not closed under intersection.*

Proof. Consider the languages $L_1 = \{a^n b^n c^m \mid m, n \geq 0\}$ and $L_2 = \{a^m b^n c^n \mid m, n \geq 0\}$. There exist DPDA recognising L_1 and L_2 . Hence by Lemma 1 there exist GFG-PDA recognising L_1 and L_2 .

Now let $L = L_1 \cap L_2 = \{a^n b^n c^n \mid n \geq 0\}$. As L is not a CFL there does not exist any GFG-PDA recognising L . \square

Lemma 9. *GFG-CFL is not closed under complementation.*

Proof. Recall the language

$$B_2 = \{a^i \$ a^j \$ b^k \$ \mid k \leq \max(i, j)\}.$$

We proved with Lemma 12 that $B_2 \in \text{GFG-CFL}$, yet Lemma 13 shows that its complement B_2^c is not even a context-free language. \square

Lemma 10. *GFG-CFL is not closed under set difference.*

Proof. Closure under set difference implies closure under complementation since for every language L over alphabet Σ , we have that the complement L^c is equal to $\Sigma^* \setminus L$. \square

Lemma 11. *GFG-CFL is not closed under homomorphism.*

Proof. The language

$$L = \left\{ \binom{a}{1}^n \binom{b}{\#}^n \mid n \geq 0 \right\} \cup \left\{ \binom{a}{2}^n \binom{b}{\#}^{2n} \mid n \geq 0 \right\}$$

is recognised by a DPDA, and hence by a GFG-PDA using Lemma 1, but its projection (which is a homomorphism)

$$\{a^n b^n \mid n \geq 0\} \cup \{a^n b^{2n} \mid n \geq 0\}$$

cannot be recognised by a GFG-PDA (see Lemma 14). \square

Closure Properties with Regular Languages

Theorem 9. *If L is in GFG-CFL and R is regular, then $L \cup R$, $L \cap R$ and $L \setminus R$ are also in GFG-CFL, but $R \setminus L$ is not necessarily in GFG-CFL.*

Proof. Consider a GFG-PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ recognising L , and a resolver r for \mathcal{P} . By definition, r only has to induce a run on every $w \in L(\mathcal{P})$, but does not necessarily induce a run on $w \notin L(\mathcal{P})$. First, we turn \mathcal{P} into an equivalent GFG-PDA \mathcal{P}' that has a resolver that induces a run on every input $w \in \Sigma^*$. This property allows us then to take the product of \mathcal{P}' and an DFA for R .

To this end, we add a fresh nonfinal sink state q_s with a self-loop (q_s, X, a, q_s, X) for every input letter $a \in \Sigma$ and every stack symbol $X \in \Gamma_\perp$. Also, we add transitions so that every configuration has, for every $a \in \Sigma$, an enabled a -transition to the sink. The resulting PDA \mathcal{P}' is equivalent and r is still a resolver for it. But, we can also turn r into a resolver r' that induces a run on every possible input as follows: If $\ell(\tau_0 \cdots \tau_n)$ is a prefix of a word in $L(\mathcal{P})$, then we define $r'(\tau_0 \cdots \tau_n, a) = r(\tau_0 \cdots \tau_n, a)$ for every $a \in \Sigma$. Otherwise, if $\ell(\tau_0 \cdots \tau_n)$ is not a prefix of a word in $L(\mathcal{P})$, we define $r'(\tau_0 \cdots \tau_n, a) = (q, X, a, q_s, X)$, where (q, X) is the mode of the last configuration of the run induced by $\tau_0 \cdots \tau_n$. Thus, as soon as the input can no longer be extended to a word in $L(\mathcal{P})$, then run induced by r' moves to the sink state and processes the remaining input.

Now, let \mathcal{A} be a DFA recognising R . For $L \cup R$, we construct the product PDA \mathcal{P}_\cup of \mathcal{P}' and \mathcal{A} that simulates a run of \mathcal{P}' and the unique run of \mathcal{A} simultaneously on an input word and accepts if either the

run of \mathcal{P} or the run of \mathcal{A} is accepting. We note that when an ε -transition is chosen in \mathcal{P}' by the resolver, then no move is made in \mathcal{A} . As \mathcal{P}' has a run on every input, the product PDA \mathcal{P}_\cup has one as well.

For $L \cap R$, we construct a PDA \mathcal{P}_\cap which is similar to \mathcal{P}_\cup with the difference that \mathcal{P}_\cap accepts when each of \mathcal{P}' and \mathcal{A} has an accepting run on the input word.

Both \mathcal{P}_\cup and \mathcal{P}_\cap are GFG-PDA since only the nondeterminism of \mathcal{P}' needs to be resolved.

Finally, since $L \setminus R = L \cap R^c$ and the complement R^c is regular, it follows that $L \setminus R$ is recognised by a GFG-PDA if there is a GFG-PDA recognising L . For $R \setminus L$, note that Σ^* is a regular language and $\Sigma^* \setminus L = L^c$ which by Lemma 9 implies that $R \setminus L$ may not be in GFG-CFL. \square

Proof of Theorem 3 We prove $\text{DCFL} \subsetneq \text{GFG-CFL} \subsetneq \text{CFL}$ in several steps. Recall that we defined $B_2 = \{a^i \$ a^j \$ b^k \$ \mid k \leq \max(i, j)\}$.

Lemma 12. $B_2 \in \text{GFG-CFL}$.

Proof. Let us summarise the behaviour of the pushdown automaton \mathcal{P}_{B_2} recognising B_2 (see Figure 2). First, the automaton copies the two blocks of a 's on the stack. Then, when it processes the second $\$$, it transitions nondeterministically to either p_1 or p_2 . In p_1 , it erases the second a -block from the stack, so that the first block is at the top of the stack, and then transitions to p_2 . In p_2 , the automaton compares the number of b 's in the input with the number of a 's in the topmost block of the stack. If the latter is larger than or equal to the former, \mathcal{P}_{B_2} pops one a for each b in the input, and then transitions to the final state when it processes the third $\$$.

When processing the second $\$$, knowing whether the first or second block of the prefix contains more a 's allows the nondeterminism to be resolved: if the first block contains more a 's, take the transition to the state p_1 , if the second block contains more a 's, take the transition to the state p_2 . \square

To show that B_2 is not in DCFL , we prove that its complement is not even context-free. This suffices, as DCFL is closed under complementation.

Lemma 13. *The complement B_2^c of B_2 is not in CFL .*

Proof. Assume, for the sake of contradiction, that the complement B_2^c of B_2 is in CFL . Now consider the regular language

$$A = \{a^i \$ a^j \$ b^k \$ \mid i, j, k \in \mathbb{N}\}.$$

Since the intersection of a context-free language and a regular language is context-free, we have that $B_2^c \cap A \in \text{CFL}$. Therefore, $B_2^c \cap A$ satisfies the pumping lemma for context-free languages: there exists $m \in \mathbb{N}$ such that the word $z = a^m \$ a^m \$ b^{m+1} \$ \in B_2^c \cap A$ can be decomposed as $z = uvwxy$ such that

1. $|vx| \geq 1$;
2. $uv^n wx^n y \in B_2^c \cap A$ for every $n \geq 0$.

Note that Item 2 directly implies that both v and x are in the language $\{a\}^* \cup \{b\}^*$, as otherwise uv^2wx^2y is not in A . On top of that, Item 1 implies that either v or x is in $\{a\}^+ \cup \{b\}^+$. We conclude by proving, through a case distinction, that Item 2 cannot hold as either uvw or uv^2wx^2y is in B_2 .

- Assume that neither v nor x is in $\{b\}^+$. Then either v or x is in $\{a\}^+$, hence $uv^2wx^2y = a^{m_1} \$ a^{m_2} \$ b^{m+1} \$$ for some $m_1, m_2 \geq m$ such that either $m_1 > m$ or $m_2 > m$. In both cases, we get $uv^2wx^2y \in B_2$.
- Assume that either v or x is in $\{b\}^+$. Then pumping v and x down in z reduces the size of at most one of the a -blocks: we have that $uvw = a^{m_1} \$ a^{m_2} \$ b^{m_3+1} \$$ for $m_1, m_2, m_3 \leq m$ such that $m_3 < m$, and either $m_1 = m$ or $m_2 = m$. In both cases, we get $uvw \in B_2$.

As every possible case results in a contradiction, our initial hypothesis is false: $B_2^c \notin \text{CFL}$. \square

The previous two lemmata and Lemma 1 yield $\text{DCFL} \subsetneq \text{GFG-CFL}$.

Now, we prove $\text{GFG-CFL} \subsetneq \text{CFL}$. Recall that we defined $L = \{a^n b^n \mid n \geq 0\} \cup \{a^n b^{2n} \mid n \geq 0\}$, which is clearly in CFL . Hence, the following lemma completes the proof of Theorem 3.

Lemma 14. $L \notin \text{GFG-CFL}$.

Proof. We show that there does not exist a GFG-PDA recognising L . In fact, we show that if there existed a GFG-PDA \mathcal{P} recognising L , then we could construct a PDA $\widehat{\mathcal{P}}$ recognising the language $\widehat{L} = L \cup \{a^n b^n c^n \mid n \geq 0\}$. Since \widehat{L} is not in CFL, we would thus reach a contradiction.

The idea behind the construction is to replicate the part of the control unit of \mathcal{P} which processes the suffix b^n of an input word $a^n b^{2n}$ with the difference that in the newly added parts, the transitions caused by input symbol b are replaced with similar ones for input symbol c . This new part of the control unit may be entered after $\widehat{\mathcal{P}}$ has processed $a^n b^n$.

We now construct the PDA $\widehat{\mathcal{P}}$ from \mathcal{P} as follows. Let $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ with $Q = \{q_0, q_1, \dots, q_n\}$, and let $q_0 = q_I$. Now consider $\widehat{\mathcal{P}} = (Q \cup \widehat{Q}, \Sigma, \Gamma, q_I, \Delta \cup \widehat{\Delta}, F \cup \widehat{F})$ with $\widehat{Q} = \{\widehat{q}_0, \widehat{q}_1, \dots, \widehat{q}_n\}$, $\widehat{F} = \{\widehat{q}_i \mid q_i \in F\}$, and $\widehat{\Delta}$ includes the following additional transitions:

1. $\{(q_f, X, \varepsilon, \widehat{q}_f, X) \mid q_f \in F, X \in \Gamma_\perp\}$: switch from the original final states to the new states.
2. $\{(\widehat{q}_i, X, c, \widehat{q}_j, \gamma) \mid (q_i, X, b, q_j, \gamma) \in \Delta\}$: replicate the original b -transitions by c -transitions in the new states.
3. $\{(\widehat{q}_i, X, \varepsilon, \widehat{q}_j, \gamma) \mid (q_i, X, \varepsilon, q_j, \gamma) \in \Delta\}$: replicate all ε -transitions.

Now we show that $L(\widehat{\mathcal{P}}) = \widehat{L}$. First we show that $L(\widehat{\mathcal{P}}) \subseteq \widehat{L}$. Consider a word $w \in L(\widehat{\mathcal{P}})$. There may be two cases:

- (i) Assume $\widehat{\mathcal{P}}$ has an accepting run on w that does not visit a state in \widehat{Q} . In this case, we have that w is in $L(\mathcal{P}) = L \subseteq \widehat{L}$.
- (ii) Assume there exists an accepting run of $\widehat{\mathcal{P}}$ on w that visits a state in \widehat{Q} . Since \mathcal{P} recognises L , and by construction of $\widehat{\mathcal{P}}$, a state $\widehat{q}_i \in \widehat{Q}$ can be reached from a state $q_i \in Q$ only if $\widehat{q}_i \in \widehat{F}$ and $q_i \in F$, and the corresponding transition is an ε -transition, we have that starting from the initial configuration (q_I, \perp) , a state in \widehat{Q} is reached for the first time only after processing an input prefix $a^n b^n$ or $a^n b^{2n}$ for some $n \geq 0$. If this prefix of w is $a^n b^{2n}$, then $w = a^n b^{2n}$. This is because if $w = a^n b^{2n} c^m$ for some $m > 0$ (recall that after visiting a state \widehat{q}_i in \widehat{Q} , the only non- ε transitions possible are on the letter c), then by the construction of $\widehat{\mathcal{P}}$, we have that \mathcal{P} can accept the word $a^n b^{2n} b^m$ which is not in the language L . On the other hand, let the prefix be $a^n b^n$ when a state $\widehat{q}_i \in \widehat{Q}$ is visited for the first time. Note that $\widehat{q}_i \in \widehat{F}$, and let $(\widehat{q}_i, \gamma_i)$ be the corresponding configuration. If a sequence of transitions $\widehat{\tau}_i \dots \widehat{\tau}_j$ from $(\widehat{q}_i, \gamma_i)$ to $(\widehat{q}_j, \gamma_j)$ is possible such that not all of $\widehat{\tau}_i \dots \widehat{\tau}_j$ are ε -transitions, that is, the transitions process c^m for some $m \in \mathbb{N}$, and $\widehat{q}_j \in \widehat{F}$, then a sequence of transitions $\tau_i \dots \tau_j$ of the same length processing b^m is possible from (q_i, γ_i) to (q_j, γ_j) with $q_j \in F$. Since this leads to an accepting run from (q_I, \perp) to (q_j, γ_j) while visiting only the states in Q on processing $a^n b^n b^m$ with $m > 0$, we have $m = n$, and hence $w = a^n b^n c^n \in \widehat{L}$.
If on the other hand, if all the transitions $\widehat{\tau}_i \dots \widehat{\tau}_j$ are ε -transitions, then $w = a^n b^n \in \widehat{L}$.

Now we prove the other direction, that is $\widehat{L} \subseteq L(\widehat{\mathcal{P}})$. Here, we rely on the fact that the accepting run of \mathcal{P} on $a^n b^n$ induced by r is a prefix of the accepting run of \mathcal{P} on $a^n b^{2n}$ induced by r . This allows to switch to the copied states \widehat{Q} after processing $a^n b^n$ and then process c^n instead of b^n .

Consider a word $w \in \widehat{L}$ such that $w \in L$. By construction of $\widehat{\mathcal{P}}$, we have that $w \in L(\widehat{\mathcal{P}})$ since $\widehat{\mathcal{P}}$ accepts all words that are also accepted by \mathcal{P} . Now suppose that $w \in \widehat{L}$ but $w \notin L$, that is, w is of the form $a^n b^n c^n$ for some $n \geq 1$. Since by assumption, we have that \mathcal{P} is a GFG-PDA recognising the language L , there exists a resolver r that for every word in L induces an accepting run of the word in L . Let (q_i, γ_i) be the configuration of \mathcal{P} reached after processing the prefix $a^n b^n$ in the run induced by r on the input $a^n b^{2n}$.

Note that $q_i \in F$ since r also induces an accepting run for the input $a^n b^n$. Now if for the input $a^n b^{2n}$, the sequence of transitions chosen by r from (q_i, γ_i) after processing $a^n b^n$ is $\tau_i, \tau_{i+1}, \dots, \tau_j$ with $(q_i, \gamma_i) \xrightarrow{\tau_i} (q_{i+1}, \gamma_{i+1}) \dots (q_{j-1}, \gamma_{j-1}) \xrightarrow{\tau_j} (q_j, \gamma_j)$, with $q_j \in F$, and the sequence τ_i, \dots, τ_j processes b^n , then by the construction of $\widehat{\mathcal{P}}$, there exists a sequence of transitions $\widehat{\tau}_i, \widehat{\tau}_{i+1}, \dots, \widehat{\tau}_j$ with $(\widehat{q}_i, \gamma_i) \xrightarrow{\widehat{\tau}_i} (\widehat{q}_{i+1}, \gamma_{i+1}) \dots (\widehat{q}_{j-1}, \gamma_{j-1}) \xrightarrow{\widehat{\tau}_j} (\widehat{q}_j, \gamma_j)$ and with $\widehat{q}_j \in \widehat{F}$ such that there is an ε -transition from (q_i, γ_i) to $(\widehat{q}_i, \gamma_i)$ and the sequence $\widehat{\tau}_i, \widehat{\tau}_{i+1}, \dots, \widehat{\tau}_j$ processes c^n , and hence $w \in L(\widehat{\mathcal{P}})$.

Thus we have that $\widehat{L} = L(\widehat{\mathcal{P}})$. Hence we show that if \mathcal{P} is a GFG-PDA, then we can construct a PDA $\widehat{\mathcal{P}}$ recognising \widehat{L} which is not a CFL, thus leading to a contradiction to our assumption that L is in GFG-CFL. \square

Proof of Theorem 4 We need to show that $B = \{a^i b^j c^k \mid i, j, k \geq 1, k \leq \max(i, j)\}$ is inherently ambiguous, i.e. for every grammar generating B there is at least one word that has two different leftmost derivations.

We use standard definitions and notation for context-free grammars as in [19]. We say that a grammar is reduced, if every variable is reachable from the start variable, every variable can be reduced to a word of terminals, and for no variable A , it holds that $A \xRightarrow{*} A$.

Let $D(G) = \{A \in V \mid A \xRightarrow{*} xAy \text{ for some } x, y \text{ with } xy \neq \varepsilon\}$. An unambiguous CFG G is called *almost-looping*, if

1. G is reduced,
2. all variables, possibly other than the start variable S , belong to $D(G)$, and
3. either $S \in D(G)$ or S occurs only once in the leftmost derivation of any word in $L(G)$.

Now we state the following lemma from [29].

Lemma 15. *For every unambiguous CFG G , there exists an unambiguous almost-looping CFG G' such that $L(G) = L(G')$.*

An important property of the grammar obtained through the translation is that if the initial grammar is unambiguous, then so is the translated almost-looping grammar.

An example of an almost-looping grammar for language B is the following:

$$\begin{array}{ll} S \rightarrow S_1 \mid S_2 & S_2 \rightarrow aS_2 \mid aD \\ S_1 \rightarrow aS_1 \mid aS_1c \mid aBc & D \rightarrow bDc \mid bD \mid bc \\ B \rightarrow bB \mid b & \end{array}$$

Fig. 4. An example CFG for language B

Now we prove the following, using techniques inspired by Maurer's proof that $\{a^i b^j c^k \mid i, j, k \geq 1, i = j \vee j = k\}$ is inherently ambiguous [29].

Lemma 16. *The language B is inherently ambiguous.*

Proof. Assume, towards a contradiction, that G is an unambiguous grammar for B , which, from Lemma 15, we can assume, without loss of generality, to be an almost-looping grammar. Let A be a variable of G .

1. A is of Type 1 if there is a derivation $A \xRightarrow{*} xAy$ where $xy = a^{n_{A,1}}$ for some $n_{A,1} > 0$.
2. A is of Type 2 if there is a derivation $A \xRightarrow{*} xAy$ where $xy = b^{n_{A,2}}$ for some $n_{A,2} > 0$.
3. A is of Type 3 if there is a derivation $A \xRightarrow{*} xAy$ where $x = a^{\ell_{A,3}}$ and $y = c^{r_{A,3}}$ for some $\ell_{A,3} \geq r_{A,3} > 0$.
4. A is of Type 4 if there is a derivation $A \xRightarrow{*} xAy$ where $x = b^{\ell_{A,4}}$ and $y = c^{r_{A,4}}$ for some $\ell_{A,4} \geq r_{A,4} > 0$.
5. A is of Type 5 if there is a derivation $A \xRightarrow{*} xAy$ where $x = a^{\ell_{A,5}}$ and $y = b^{r_{A,5}}$ for some $\ell_{A,5}, r_{A,5} > 0$.

Note that some variables may be of multiple types (e.g. the variable D in Figure 4 has Type 2 and Type 4). First, we show that each variable in $D(G)$ has at least one of these five types. So, let $A \in D(G)$. Then, there exists a derivation $A \xRightarrow{*} xAy$ with $xy \neq \varepsilon$. Note that both x and y belong to a^* , b^* , or c^* since otherwise, due to G being reduced, one could derive words that are not in the language. Next, we note that the cases where x belongs to c^* , and y belongs to a^* or b^* cannot happen. Similarly, the case where x belongs to b^* , and y belongs to a^* cannot happen. Also we cannot have xy in c^* , since this will allow us to have words with arbitrary number of c 's which can be more than the number of a 's and b 's and such a word is not in the language.

Further, we cannot have $x = a^\ell$ and $y = c^r$ with $0 < \ell < r$. Otherwise, consider a derivation of some word in B that uses A , i.e.

$$S \xRightarrow{*} \alpha A \beta \xRightarrow{*} a^s b^u c^v \quad \text{with } v \leq \max(s, u).$$

Now, towards a contradiction assume we indeed have

$$A \xrightarrow{*} xAy \quad \text{with } x = a^\ell, y = c^r \text{ and } \ell < r.$$

Then, pumping q copies of x and y , for some suitable $q \in \mathbb{N}$, yields a derivation

$$S \xrightarrow{*} \alpha A \beta \xrightarrow{*} \alpha x^q A y^q \beta \xrightarrow{*} a^{s+\ell q} b^u c^{v+r q} \quad \text{such that } v + r q > \max(s + \ell q, u),$$

i.e. we have derived a word that is not in B . Similarly, we cannot have $x = b^\ell$ and $y = c^r$ for some $0 < \ell < r$. Altogether, this implies that A indeed has at least one of the five types stated above.

Moreover, we claim that there is a $t \in \mathbb{N}$ such that the following three properties are true for every word $w \in B$:

Property 1 If w has more than t c 's, then the (unique) leftmost derivation of w has the form

$$S \xrightarrow{*} \alpha A \beta \xrightarrow{*} \alpha x A y \beta \xrightarrow{*} w \quad \text{such that } xy \text{ contains a } c.$$

Thus, A has type 3 or type 4.

Property 2 If w has more than t a 's, then the (unique) leftmost derivation of w has the form

$$S \xrightarrow{*} \alpha A \beta \xrightarrow{*} \alpha x A y \beta \xrightarrow{*} w \quad \text{such that } xy \text{ contains an } a.$$

Thus, A has type 1, type 3, or type 5.

Property 3 If w has more than t b 's, then the (unique) leftmost derivation of w has the form

$$S \xrightarrow{*} \alpha A \beta \xrightarrow{*} \alpha x A y \beta \xrightarrow{*} w \quad \text{such that } xy \text{ contains a } b.$$

Thus, A has type 2, type 4, or type 5.

We prove these properties as follows: we denote by d the *width* of the grammar G which is the maximum number of symbols appearing on the right side of some production rule of G . Further, we denote by m the number of variables appearing in G . We argue that $t = d^{m+1}$ satisfies the three properties above. We focus on Property 1, the two other proofs are similar. Suppose that w contains more than d^{m+1} c 's and consider the derivation tree of that word. The *weight* $\omega(v)$ of a vertex v in the derivation tree is defined as the number of c 's in the subtree rooted at v . Hence, the root of the derivation tree has at least weight d^{m+1} . We build a finite path v_0, v_1, \dots, v_k from the root of this tree to one of its leaves as follows: The initial vertex v_0 is the root and at each step, we choose as successor of v_i its child v_{i+1} with the largest weight. A vertex v_i of this path is *decreasing* if $\omega(v_i) > \omega(v_{i+1})$. There are at least $m + 1$ decreasing vertices on the path because $\omega(v_0) = d^{m+1}$, $\omega(v_k) = 1$, and $\omega(v_{i+1}) \geq \frac{1}{d} \cdot \omega(v_i)$. Thus, there are two decreasing vertices on the path that are labeled by the same variable A such that there is a derivation of the form $A \xrightarrow{*} xAy$ with some c in xy .

Let $p > t$ be a positive integer divisible by the least common multiple of the $n_{A,i}, \ell_{A,i}$ and $r_{A,i}$ for all $A \in D(G)$ and $i \in \{1, \dots, 5\}$, where we define $n_{A,1} = 1$ if A is not of Type 1, and similarly for all other $i > 1$. We show that the word $w = a^{2p} b^{2p} c^{2p} \in B$ has two leftmost derivations.

First consider the derivation of the word $w_b = a^{2p} b^p c^{2p} \in B$. As we have more than t c 's in w_b Property 1 shows that the derivation contains a variable of Type 3 or Type 4. Next, we argue that it cannot contain a variable of Type 4: The occurrence of such a variable would allow us to either produce a word that is not in $a^* b^* c^*$ or to inject $b^p c^r$ for $p \geq r > 0$ leading to the derivation of $a^{2p} b^{2p} c^{2p+r}$, which is not in the language. Thus, the derivation of w_b uses at least one variable of Type 3. Also, since w_b has $p > t$ b 's, Property 3 implies that the (unique) leftmost derivation of w_b has the form

$$S \xrightarrow{*} \alpha A \beta \xrightarrow{*} \alpha x A y \beta \xrightarrow{*} w_b \quad \text{such that } xy \text{ contains a } b.$$

Thus A is a variable of Type 2 or Type 5 (note that we have already ruled out Type 4 above). More precisely, we have that x belongs to a^+ or b^* and $y = b^j$ for some $j \in \mathbb{N}$. Now we show that the case where x belongs to a^+ is not possible. Assume for contradiction that $x = a^i$ for some $i > 0$. Then we also have the derivation

$$S \xrightarrow{*} \alpha A \beta \xrightarrow{*} a^{2p-i} b^{p-j} c^{2p} \notin B.$$

Therefore, A is a Type 2 variable that is used in the derivation of w_b , which can be used to inject another b^p , yielding a derivation of w . Thus, we have exhibited a derivation of w that uses a variable of Type 3.

Now consider a derivation of the word $w_a = a^p b^{2p} c^{2p}$. Such a derivation cannot contain a variable of Type 3 since this allows us either to produce a word that is not in $a^* b^* c^*$ or to inject $a^p c^r$ for $p \geq r > 0$, leading to the derivation of $a^{2p} b^{2p} c^{2p+r} \notin B$. Further, arguing as above, some variable of Type 1 must appear in the derivation of w_a that is used to obtain sufficient number of a 's in the derivation of w_a . Such a variable of Type 1 can be used to inject a^p into w_a which leads to the derivation of w . Thus, we have exhibited a derivation of w that does not contain a variable of Type 3.

Altogether, there are two different leftmost derivations of the word w . Thus, G is not unambiguous, yielding the desired contradiction. \square

Proof of Theorem 5 Recall that we need to prove that GFG-PDA can be exponentially more succinct than DPDA, and that PDA can be double-exponentially more succinct than GFG-PDA.

We first consider the gap between DPDA and GFG-PDA. Recall that we defined $c_n \in (\{0, 1\}^n)^*$ to be the word describing an n -bit binary counter counting from 0 to $2^n - 1$ and

$$C_n = \{w \in \{0, 1, \$, \#\}^* \mid w \neq c_n \#\}.$$

We prove that C_n is recognised by a PDA of linear size, but every GFG-PDA recognising C_n has exponential size. The proof is split into two parts.

Lemma 17. *The language C_n is recognised by a GFG-PDA of size $O(n)$.*

Proof. We define a PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ that recognises C_n . The automaton \mathcal{P} operates in three phases: a push phase, followed by a check phase, and then a final phase. These phases work as follows. Suppose that \mathcal{P} receives an input $w \in \{0, 1, \$, \#\}^*$. During the first phase, \mathcal{P} pushes the input processed onto the stack until the sequence 1^n appears. If it never appears, the input is accepted. During this phase, \mathcal{P} also checks whether the prefix w' of w processed up to this point is a sequence of counter values starting with 0^n , i.e. whether w' is in the language

$$L_c = \{\$d_0\$d_1\$ \dots \$d_m \mid d_0 = 0^n, d_i = 1^n \Leftrightarrow i = m, \text{ and } d_i \in \{0, 1\}^n \text{ for all } 1 \leq i \leq m\}.$$

If $w' \notin L_c$, then \mathcal{P} immediately accepts. Otherwise, \mathcal{P} moves to the second phase. During the check phase, \mathcal{P} pops the stack. At any point, \mathcal{P} can nondeterministically guess that the top symbol of the stack is evidence of bad counting. It then accepts the input if the guess was correct. If \mathcal{P} completely pops the stack without correctly guessing an error in the counter, it moves to the final phase. Since the prefix w' processed up to this point ends with the sequence 1^n , if \mathcal{P} now processes any suffix different from a single $\#$, then the input is not equal to $c_n \#$, and can be accepted.

The stack alphabet of \mathcal{P} has constant size 3. The push phase requires $3(n+1)$ states:

- First, \mathcal{P} checks whether $\$0^n$ is a prefix of the input. This can be done with $n+2$ states.
- Then, \mathcal{P} checks whether the following $\{0, 1\}^*$ segments are n -bits wide, and only the last one is 1^n . This can be done with $2n+1$ additional states: repeatedly, \mathcal{P} processes $n+1$ symbols, checks whether only the first of them is a $\$$, and keeps track of whether at least one of them is 0.

We now show that $6(n+1)$ additional states are enough for the check phase. To this end, we study the errors that \mathcal{P} needs to check. Note that, to increment the counter correctly, we need to change the value of all the bits starting from the last 0, and leave the previous bits unchanged. Therefore, \mathcal{P} can recognise with $6(n+1)$ states whether the top symbol of the stack does not correspond to a correct counter increment: \mathcal{P} pops the top $n+1$ stack symbols while keeping in memory

- the value of the first symbol popped;
- whether we have not yet popped a $\$$ (there is exactly one $\$$ in the top $n+1$ stack symbols, as the stack content is in L_c), or a $\$$ but no 0 afterwards, or a $\$$ and at least one 0 afterwards.

The input is accepted whenever the first symbol popped and the top stack symbol after popping match yet no 0 has been popped between the $\$$ and the last symbol, or they differ yet at least one 0 has been popped between the $\$$ and the last symbol.

Finally, only three states are needed for the final phase: when the bottom of the stack is reached, \mathcal{P} transitions to a new state, and from there it checks whether the suffix is in the language $\{0, 1, \$, \#\}^* \setminus \{\#\}$.

To conclude, note that \mathcal{P} is good-for-games: the only nondeterministic choice happens during the check phase, and the resolver knows which symbols of the stack are evidence of bad counting. Note that this choice only depends on the current stack content. \square

Lemma 18. *Every DPDA recognising the language C_n has at least exponential size in n .*

Proof. It is known that every DPDA can be complemented at the cost of multiplying its number of states by three [19]. Therefore, to prove the statement, we show that even every PDA recognising the complement $\{c_n\#\}$ of C_n has at least exponential size in n :

Claim. Every PDA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ recognising $\{c_n\#\}$ has a size greater than $2^{(n-1)/3}$.

To prove the claim, we transform \mathcal{P} into a context-free grammar generating the singleton language $\{c_n\#\}$, and then we show that such a grammar requires exponentially many variables. This is a direct consequence of the *mk* Lemma [10], but proving it directly using similar techniques yields a slightly better bound.

Before changing \mathcal{P} into a grammar, we slightly modify its acceptance condition: we add to \mathcal{P} a fresh final state f in which the stack can be completely popped *including the bottom of stack symbol* \perp (which normally cannot be touched according to our definition of PDA). Moreover, we allow \mathcal{P} to transition towards f nondeterministically from all of its other final states. This new automaton, which accepts by empty stack, is easily transformed into a grammar \mathcal{G} using the standard transformation [19]:

- The terminals of \mathcal{G} are $0, 1, \$$ and $\#$.
- The variables of \mathcal{G} are the triples (p, X, q) , for every state $p, q \in Q \cup \{f\}$ and stack symbol $X \in \Gamma_{\perp}$.
- The initial variable is (q_I, \perp, f) , where q_I is the initial state of \mathcal{P} and f is the fresh final state.
- Each transition $(p, X, a, q, \gamma) \in \Delta$ yields production rules as follows:
 1. If $\gamma = \varepsilon$, then \mathcal{G} has the production rule $(p, X, q) \rightarrow a$;
 2. If $\gamma = Y$, then \mathcal{G} has the production rule $(p, X, q_1) \rightarrow a(q, Y, q_1)$ for all $q_1 \in Q$;
 3. If $\gamma = YZ$, then \mathcal{G} has the production rule $(p, X, q_2) \rightarrow a(q, Y, q_1)(q_1, Z, q_2)$ for all $q_1, q_2 \in Q$.

The variables can be interpreted as follows: for every $p, q \in Q$ and $X \in \Gamma$, the variable (p, X, q) can be derived into any input word $w \in \{0, 1, \$\}^*$ that \mathcal{P} can process starting in state p and ending in state q while consuming the symbol X from the top of the stack. Therefore, in particular, since the initial variable is (q_I, \perp, f) , \mathcal{G} generates the same language as \mathcal{P} .

We now prove that the grammar \mathcal{G} has at least 2^{n-1} distinct variables, hence $(|Q|+1)^2(|\Gamma|+1) \geq 2^{n-1}$, which implies that the size $|Q| + |\Gamma|$ of \mathcal{P} is at least $\geq 2^{(n-1)/3}$. To this end, we study a (directed) derivation tree T of the word $c_n\#$.

Remember that $c_n = \$d_0\$d_1\$ \dots \d_{2^n-1} represents an n -bit binary counter counting from 0 to $2^n - 1$. For each $0 \leq i \leq 2^n - 1$, let us consider the vertex v_i of T such that the counter value d_i is an infix of the derivation of v_i , but of none of its children. In other words, d_i is split between the derivations of the children of v_i . By definition of the grammar \mathcal{G} , each vertex of T has at most three children, hence at most two counter values can be split amongst the children of a given vertex, which implies that $v_i \neq v_{i+2}$ for all $0 \leq i \leq 2^{n-3}$. Therefore, the vertices $v_0, v_2, v_4, \dots, v_{2^n-2}$ are all distinct. Finally, since $c_n\#$ is the only word recognised by \mathcal{P} and each counter value $d \in \{0, 1\}^n$ appears a single time as an infix of $c_n\#$, the 2^{n-1} variables labelling these vertices need to be distinct. \square

Now, we consider the gap between GFG-PDA and PDA. Recall that L_n is the language of words over $\{0, 1\}$ such that the n^{th} bit from the end is a 1. We need to show that there exists a PDA of size $O(\log n)$ recognising L_n , and that every GFG-PDA recognizing L_n has exponential size. Again, the proof is split into two parts.

Lemma 19. *There exists a PDA of size $O(\log n)$ recognising L_n .*

Proof. We describe a PDA \mathcal{P} that recognises L_n . The PDA \mathcal{P} nondeterministically guesses the n^{th} bit from the end, checks that it is a 1 and switches to a counting gadget that checks that the word ends in n steps, as follows:

- (i) It pushes the binary representation of $n - 2$ onto the stack. For example, if $n = 8$, then 110 is pushed onto the stack with 0 at the top. Note that $\log(n - 2)$ states suffice for pushing the binary representation of $n - 2$. If $n = 1$, then instead of pushing anything onto the stack, the automaton directly moves to a final state without any enabled transitions.
- (ii) Then \mathcal{P} moves to a state that attempts to decrement the counter by one for each successive input letter, as follows: When an input letter is processed, it pops m 0's until 1 is at the top of the stack, replaces the 1 with a 0, and finally pushes m 1's back onto the stack before processing the next letter. If the stack empties before a 1 is at the top of the stack, then the counter value is 0 and the automaton moves to a final state with no enabled transitions. Note that $O(\log n)$ states again suffice for this step.

Thus \mathcal{P} has $O(\log n)$ states. Note that for all n , \mathcal{P} uses only three stack symbols that are 0, 1, and \perp . Thus the size of \mathcal{P} is $O(\log n)$, and \mathcal{P} recognises L_n . \square

Lemma 20. *Every GFG-PDA recognising L_n has at least exponential size in n .*

Towards proving this, we define the following notions. We say that a word w of length n is *rotationally equivalent* to a word w' if w' is obtained from w by rotating it. For example, the word $w = 1101$ is rotationally equivalent to $w' = 1110$ since w' can be obtained from w by rotating it once to the right. Note that the words that are rotationally equivalent form an equivalence class, and thus rotational equivalence partitions $\{0, 1\}^n$. Since the size of each class is at most n , the number of equivalence classes is at least $\frac{2^n}{n}$.

Now, we define the *stack height* of a configuration $c = (q, \gamma)$ as $\text{sh}(c) = |\gamma| - 1$, and we define *steps* of a run as usual: Consider a run $c_0\tau_0c_1\tau_1 \cdots c_{n-1}\tau_{n-1}c_n$. A position s is a step if for all $s' \geq s$, we have that $\text{sh}(c_{s'}) \geq \text{sh}(c_s)$, that is, the stack height is always at least $\text{sh}(c_s)$ after position s . Any infinite run of a PDA has infinitely many steps. We have the following observation.

Proposition 1. *If two runs of a PDA have steps s_0 and s_1 , respectively, with the same mode, then the suffix of the run following the step s_0 can replace the suffix of the other run following the step s_1 , and the resultant run is a valid run of the PDA.*

Now, we are ready to prove Lemma 20. Here, we work with infinite inputs for GFG-PDA. The run induced by a resolver on such an input is the limit of the runs on the prefixes.

Proof. Let \mathcal{P} be a GFG-PDA with resolver r that recognises L_n with a set Q of states and a stack alphabet Γ . We show that $|Q| \cdot |\Gamma| \geq \frac{2^n}{n}$.

Towards a contradiction, assume that $|Q| \cdot |\Gamma| < \frac{2^n}{n}$. Then there exist two words w_0 and w_1 of length n that are not rotationally equivalent and such that the runs ρ_0 and ρ_1 of \mathcal{P} induced by r on w_0^ω and w_1^ω contain steps with the same mode, at positions s_0 and s_1 in ρ_0 and ρ_1 respectively, such that at least n letters are processed before s_0 and s_1 . Now consider in each of these two runs the sequence of input letters of length n preceding and including the step position. Let these n letter words be w'_0 and w'_1 respectively. Since w_0 and w_1 are not rotationally equivalent, w'_0 and w'_1 differ in at least one position $j \leq n$.

W.l.o.g., assume that for w'_0 , the bit at position j is 0, while it is 1 at position j for w'_1 . Since the resolver chooses a run such that for every word where the n^{th} letter from the end is a 1 is accepted, this implies that ρ_0 does not visit a final state after processing $j - 1$ letters after s_0 , while ρ_1 visits a final state after processing $j - 1$ letters after s_1 .

Now we reach a contradiction as follows. The suffix of ρ_0 starting from position $s_0 + 1$ can be replaced with the suffix of ρ_1 starting from position $s_1 + 1$. By Proposition 1, this yields a valid run ρ of \mathcal{P} . However, since the state that occurs after $j - 1$ letters are processed after position s_1 in ρ_1 is final, after the replacement, the state that occurs after $j - 1$ letters are processed after position s_0 in ρ is final as well. However the n^{th} letter from the end of the word processed by this accepting run of \mathcal{P} is a 0, contradicting that \mathcal{P} recognises L_n . Thus we have that $|Q| \cdot |\Gamma|$ is at least equal to the number of rotationally equivalent classes, that is, $|Q| \cdot |\Gamma| \geq \frac{2^n}{n}$. Thus the size of \mathcal{P} is at least $(\frac{2^n}{n})^{1/2}$. \square

Proof of Theorem 6 Recall that we need to prove that GFG-VPA can be exponentially more succinct than DVPA and that VPA can be exponentially more succinct than GFG-VPA. We split the proof into two parts.

Lemma 21. *GFG-VPA can be exponentially more succinct than DVPA.*

Proof. We construct a family $(C'_n)_{n \in \mathbb{N}}$ of languages such that there is a GFG-VPA of size $O(n)$ recognising C'_n , yet every DVPA recognising C'_n has at least exponential size in n . This family is obtained by adapting the family $(C_n)_{n \in \mathbb{N}}$ that we used to prove the succinctness of GFG-PDA in Section 4: Once again, we consider the word $c_n \in (\{0, 1\}^n)^*$ describing an n -bit binary counter counting from 0 to $2^n - 1$. We consider the languages $C'_n \subseteq \{0, 1, \$, \#\}^*$ of bad counters, where 0, 1 and \$ are call symbols and # is a return symbol:

$$C'_n = \{w \in \{0, 1, \$, \#\}^* \mid w \neq c_n \#^{2^n(n+1)}\}$$

The only difference with C_n is that the forbidden word is $c_n \#^{2^n(n+1)}$ instead of $c_n \#$. A GFG-VPA of size $O(n)$ recognising C'_n is obtained by a small modification of the construction presented in the proof of Lemma 17. We adapt the construction of the automaton \mathcal{P} recognising C_n as follows:

- The push phase is identical;
- The check phase is performed by consuming the # symbols instead of having ε -transitions. While the stack is not empty, \mathcal{P} accepts even if it has not found evidence of bad counting yet. Moreover, \mathcal{P} transitions towards a final sink state if a non-# symbol is read. Once the stack is empty, it transitions towards the final phase;
- In the final phase, since the prefix processed up to this point ends with an empty stack, if the suffix left to read is non-empty then the input is not equal to $c_n \#^{2^n(n+1)}$, and can be accepted.

Finally, we can prove that every DPDA (and in particular every DVPA) recognising C'_n has at least exponential size in n in the exact same way as we proved Lemma 18: The functions term and chunk used to prove the statement ignore the # symbols, hence C_n and C'_n can be treated identically. Note that this lower bound is independent of the partition of the letters into calls, returns, and internals. \square

Lemma 22. *VPA can be exponentially more succinct than GFG-VPA.*

We show that there exists a family $(L'_n)_{n \in \mathbb{N}}$ of languages such that there exists a VPA of size $O(n)$ recognising L'_n while every GFG-VPA recognising the same language has size at least $2^{n/6}$.

Towards this we consider a language L'_n of words in $(01 + 10)^* \cdot (\varepsilon + 0 + 1)$ with the n^{th} last letter being 1. We first note that L'_n can be recognised by a VPA with $O(n)$ states, which checks that the input is in $(01 + 10)^* \cdot (\varepsilon + 0 + 1)$ and nondeterministically guesses the n^{th} last letter and verifies that it is a 1.

First, we claim that every DFA recognising L'_n has exponential size.

Remark 3. Every DFA recognising L'_n has at least $2^{\lceil n/2 \rceil}$ states.

Using this, we obtain an exponential lower bound on the size of GFG-VPA recognising L'_n , thereby completing the proof of Lemma 22.

Lemma 23. *Every GFG-VPA recognising L'_n has at least size $2^{\lceil n/6 \rceil}$.*

Proof. The proof is based on the fact that GFG-NFA can be determinised by pruning [6], that is, they always contain an equivalent DFA, i.e. the lower bound of Remark 3 is applicable to GFG-NFA as well.

Let \mathcal{P} be a GFG-VPA recognising L'_n . We consider the following cases:

1. *Both 0 and 1 are either a return symbol or an internal symbol:* The GFG-VPA \mathcal{P} in this case can essentially be seen as a GFG-NFA with the same set of states, since the stack is not used (it is always equal to \perp). Given that GFG-NFA are determinisable by pruning, by Remark 3 such a GFG-NFA has at least $2^{\lceil n/2 \rceil}$ states.
2. *At least one of 0 and 1 is a call symbol while the other one is a call or an internal symbol:* Let Q be the set of states and Γ be the stack alphabet of \mathcal{P} . Since the height of the stack is nondecreasing, \mathcal{P} has only access to the top stack symbol. We can thus construct an equivalent GFG-NFA over finite words with states in $Q \times \Gamma$. Since GFG-NFA are determinizable by pruning, and using Remark 3 again, we have that $|Q| \cdot |\Gamma| \geq 2^{\lceil n/2 \rceil}$. Thus either $|Q| \geq 2^{n/4}$ or $|\Gamma| \geq 2^{n/4}$. Hence for this case, we have that the size of the GFG-VPA is at least $2^{n/4}$.
3. *One of 0 and 1 is a call symbol while the other one is a return symbol:* Note that since a word in L'_n is composed of sequences of 10 and 01, the stack height can always be restricted to 2. Thus the configuration space of \mathcal{P} , restricted to configurations on accepting runs, is finite, and there is an equivalent GFG-NFA of size at most $|Q| \cdot |\Gamma|^2$. Thus $|Q| \cdot |\Gamma|^2 \geq 2^{\lceil n/2 \rceil}$ giving either $|Q| \geq 2^{\lceil n/6 \rceil}$ or $|\Gamma| \geq 2^{\lceil n/6 \rceil}$. Again by the determinizability by pruning argument, we have that the size of the GFG-VPA \mathcal{P} is at least $2^{\lceil n/6 \rceil}$. \square

Proof of Theorem 7 We need to prove that GFGness of VPA is decidable in EXPTIME.

We use the *one-token game*, introduced by Bagnol and Kuperberg [3] in the context of regular languages. Given a VPA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$, the positions of the one-token game consist of pairs of configurations (c_i, c'_i) , starting from initial configuration of \mathcal{P} . At each round i :

- Player 1 picks a letter $a_i \in \Sigma$,
- Player 2 picks an a_i -transition $\tau_i \in \Delta$ enabled in c_i , leading to a configuration c_{i+1} ,
- Player 1 picks an a_i -transition $\tau'_i \in \Delta$ enabled in c'_i , leading to a configuration c'_{i+1} ,
- The game proceeds from the configuration (c_{i+1}, c'_{i+1}) .

A play consists of an infinite word $a_0 a_1 \dots \in \Sigma^\omega$ and two sequences of transitions $\tau_0 \tau_1 \dots$ and $\tau'_0 \tau'_1 \dots$ built by Players 2 and 1 respectively. Player 1 wins if for some n , $\tau'_0 \dots \tau'_n$ is an accepting run of \mathcal{P} over $a_0 \dots a_n$ and $\tau_0 \dots \tau_n$ is not. Recall that VPA don't have ε -transitions, so the two runs proceed in lockstep.

Observe that this game can be seen as a safety game on a visibly pushdown arena and can therefore be encoded as a Gale-Stewart game with a DCFL winning condition. This in turn is solvable in EXPTIME [40]. We now argue that this game characterises whether the VPA \mathcal{P} is GFG.

Proof. We now argue that \mathcal{P} is GFG if and only if Player 2 wins the one-token game on \mathcal{P} . One direction is immediate: if \mathcal{P} is GFG, then the resolver is also a strategy for Player 2 in the one-token game.

For the converse direction, consider the family of *copycat strategies* for Player 1 that copy the transition chosen by Player 2 until she plays an a -transition from a configuration c to a configuration c' such that there is a word aw that is accepted from c but w is not accepted from c' . We call such transitions non-residual. If Player 2 plays such a non-residual transition, then the copycat strategies stop copying and instead play the letters of w and the transitions of an accepting run over aw from c .

If Player 2 wins the one-token game with a strategy s , she wins, in particular, against this family of copycat strategies for Player 1. Observe that copycat strategies win any play along which Player 2 plays a non-residual transition. Therefore s must avoid ever playing a non-residual transition. We can now use s to induce a resolver r_s for \mathcal{P} : r_s maps a sequence of transitions over a word w to the transition chosen by s in the one-token game where Player 1 played w and a copycat strategy. Then, r_s never produces a non-residual transition. As a result, if a word w is in $L(\mathcal{P})$, then the run induced by r_s over every prefix v of w leads to a configuration that accepts the remainder of w . This is in particular the case for w itself, for which r_s induces an accepting run. This concludes our argument that r_s is indeed a resolver, and \mathcal{P} is therefore GFG.

Thus, to decide whether a VPA \mathcal{P} is GFG it suffices to solve the one-token game on \mathcal{P} , which can be done in exponential time. \square

Proof of Corollary 2 We prove that the GE-synthesis problem for GFG-VPA and DVPA is as hard as the GFGness problem for VPA. Note that this is a more general reduction that we use here only for the VPA case.

Proof. We first reduce the good-enough synthesis problem to the GFGness problem. Given a GFG-VPA $\mathcal{P} = (Q, \Sigma_1 \times \Sigma_2, \Gamma, q_I, \Delta, F)$, with resolver r , let \mathcal{P}' be \mathcal{P} projected onto the first component: $\mathcal{P}' = (Q, \Sigma_1, \Gamma, q_I, \Delta', F)$ has the same states, stack alphabet and final states as \mathcal{P} , but has an a -transition for some $a \in \Sigma_1$ whenever \mathcal{P} has the same transition over (a, b) for some $b \in \Sigma_2$. Let each transition of \mathcal{P}' be annotated with the Σ_2 -letter of the corresponding \mathcal{P} -transition. Thus \mathcal{P}' recognises the projection of $L(\mathcal{P})$ on the first component.

A resolver for \mathcal{P}' induces a GE-synthesis function for \mathcal{P} by reading off the Σ_2 -annotation of the chosen transitions in \mathcal{P}' . Indeed, the resolver produces an accepting run with annotation w' of \mathcal{P}' for every word w in the projection of $L(\mathcal{P})$ on the first component. The same run is an accepting run in \mathcal{P} over (w, w') which is therefore in $L(\mathcal{P})$. Conversely a GE-synthesis function f for \mathcal{P} , combined with r , induces a resolver r' for \mathcal{P}' by using f to choose output letters and r to choose which transition of \mathcal{P} to use; together these uniquely determine a transition in \mathcal{P}' . Then, if $w \in L(\mathcal{P}')$, f guarantees that the annotation of the run induced by r' in \mathcal{P}' is a witness w' such that $(w, w') \in \mathcal{P}$, and then r guarantees that the run is accepting, since the corresponding run in \mathcal{P} over (w, w') must be accepting.

We now reduce the GFGness problem of a VPA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ to the GE-synthesis problem of a DVPA $\mathcal{P}' = (Q, \Sigma \times \Delta, \Gamma, q_I, \Delta', F)$. The deterministic automaton \mathcal{P}' is as \mathcal{P} except that each

transition τ over a letter a in Δ is replaced with the same transition over (a, τ) in Δ' . In other words, \mathcal{P}' recognises the accepting runs of \mathcal{P} and its GE-synthesis problem asks whether there is a function that constructs on-the-fly an accepting run for every word in $L(\mathcal{P})$, that is, whether \mathcal{P} has a resolver. \square

Proof of Lemma 3 Recall that we need to prove that every GFG-PDA has a positional resolver.

Proof. Let r' be a (not necessarily positional) resolver for \mathcal{P} . We define a resolver r such that for each configuration and input letter, it makes a choice consistent with r' for some input leading to this configuration. In other words, for every reachable configuration c , let ρ_c be an input to r' inducing a run ending in c . Then, we define $r(\rho, a) = r(\rho_c, a)$, where c is the last configuration of the run induced by ρ .

We claim that r , which is positional by definition, is a resolver. Towards a contradiction, assume that this is not the case, i.e. there is a word $w \in L(\mathcal{P})$ such that the run ρ induced by r is rejecting. Since this run is finite and $w \in L(\mathcal{P})$, there is some last configuration c along the run ρ from which the rest of the word, say u , is accepted⁴ (by some other run of \mathcal{P} having the same prefix as ρ up to configuration c). Let τ be the next transition along ρ from c . Since r chose τ , the resolver r' also chooses τ after some history leading to c , over some word v . Since u is accepted from c , the word vu is in $L(\mathcal{P})$; since r' is a resolver, there is an accepting run over u from c starting with τ , contradicting that c is the last position on ρ from where the rest of the word could be accepted. \square

Proof of Lemma 4 We need to prove that the GFG-PDA \mathcal{P}_{B_3} defined in Section 6 has no pushdown resolver.

Proof. Towards a contradiction, assume that there is a pushdown resolver r for \mathcal{P}_{B_3} , implemented by a PDT $\mathcal{T} = (\mathcal{D}, \lambda)$.

From \mathcal{T} , for each $i \in \{1, 2, 3\}$, we can construct a PDA \mathcal{D}_i that recognises the language of words $w \in a^+ \$ a^+ \$ a^+$ such that \mathcal{T} chooses from q_3 the transition of \mathcal{P}_{B_3} going to p_i when constructing a run on $w \$$: this is simply the pushdown automaton \mathcal{D} underlying \mathcal{T} where inputs (transitions of \mathcal{P}_{B_3}) are projected onto their input letter in $\{a, b, \$\}$ and states q of \mathcal{T} such that $\lambda(q)(q_3, X, \$) = (q_3, \$, X, p_i, X)$ are made final, intersected with a DFA checking that the input is in $a^+ \$ a^+ \$ a^+$.

Since \mathcal{T} implements a resolver for \mathcal{P} , each \mathcal{D}_i only accepts words of the form $a^{m_1} \$ a^{m_2} \$ a^{m_3}$ such that $\max(m_1, m_2, m_3) = m_i$. Furthermore, at least for one $i \in \{1, 2, 3\}$, \mathcal{D}_i accepts $a^m \$ a^m \$ a^m$ for infinitely many m .

To reach a contradiction, we now argue that this \mathcal{D}_i recognises a language that is not context-free. Indeed, if it were, then by applying the pumping lemma for context-free languages, there would be a large enough m such that the word $a^m \$ a^m \$ a^m \in L(\mathcal{D}_i)$ could be decomposed as $uvwyz$ such that $|vy| \geq 1$ and $uv^n wy^n z$ is in the language of \mathcal{D}_i for all $n \geq 0$. In this decomposition, v and y must be $\$$ -free. Then, if either v or y occurs in the i th block and is non-empty, by setting $n = 0$ we obtain a contradiction as the i th block is no longer the longest. Otherwise, we obtain a similar contradiction by setting $n = 2$. In either case, this shows that \mathcal{T} is not a pushdown resolver for \mathcal{P} . \square

Proof of Theorem 8 We need to prove that every GFG-VPA has a (visibly) pushdown resolver.

Proof. Fix a VPA $\mathcal{P} = (Q, \Sigma, \Gamma, q_I, \Delta, F)$ and consider the following two-player game $\mathcal{G}(\mathcal{P})$, introduced by Henzinger and Piterman to decide GFGness of ω -automata [16]. In each round, first Player 1 picks a letter from Σ or ends the play. If he has not ended the play, then Player 2 picks a transition of \mathcal{P} . Hence, once Player 1 has stopped the play, Player 1 has picked an input word w over Σ^* and Player 2 has indicated a run ρ of \mathcal{P} . A finite play with outcome (w, ρ) is winning for Player 2 if either $w \notin L(\mathcal{P})$ or ρ induces an accepting run of \mathcal{P} on w .

A strategy for Player 2 in this game is a mapping $\sigma: \Sigma^+ \rightarrow \Delta$ and an outcome $(w(0) \cdots w(k), \rho(0) \cdots \rho(k))$ is consistent with σ , if $\rho(j) = \sigma(w(0) \cdots w(j))$ for every $0 \leq j \leq k$. We say that σ is winning for Player 2, if every outcome of a finite play that is consistent with σ is winning for her (note that we disregard infinite plays).

⁴ Observe that this is no longer true over infinite words as an infinite run can stay within configurations from where an accepting run exists without being itself accepting. In fact, the lemma does not even hold for coBüchi automata [23] as the existence of positional resolvers implies determinisability by pruning.

Now, Player 2 wins $\mathcal{G}(\mathcal{P})$ if and only if \mathcal{P} is a GFG-VPA. This follows as every winning strategy for Player 2 can be turned into a resolver and vice versa.

Now, as the class of languages recognized by VPA, is closed under complementation and union [2], one can encode $\mathcal{G}(\mathcal{P})$ as a Gale-Stewart game with a VPL winning condition. Such games can be solved effectively [27] and the winner always has a winning strategy implemented by a (visibly) PDT. Thus, if \mathcal{P} is a GFG-VPA, i.e. Player 2 wins $\mathcal{G}(\mathcal{P})$, then she has a winning strategy implemented by a (visibly) PDT, which can easily be turned into a (visibly) pushdown resolver for \mathcal{P} . \square