

Subcubic Certificates for CFL Reachability

Dmitry Chistikov¹ Rupak Majumdar² Philipp Schepper³

¹Centre for Discrete Mathematics and its Applications (DIMAP) & Department of
Computer Science, University of Warwick, Coventry, United Kingdom
`d.chistikov@warwick.ac.uk`

²Max Planck Institute for Software Systems, Kaiserslautern, Germany
`rupak@mpi-sws.org`

³ CISPA Helmholtz Center for Information Security, Saarbrücken, Germany &
Saarbrücken Graduate School of Computer Science, Saarland Informatics Campus,
Germany
`philipp.schepper@cispa.saarland`

Abstract

Many problems in interprocedural program analysis can be modeled as the context-free language (CFL) reachability problem on graphs and can be solved in cubic time. Despite years of efforts, there are no known truly sub-cubic algorithms for this problem. We study the related *certification* task: given an instance of CFL reachability, are there small and efficiently checkable certificates for the existence and for the non-existence of a path? We show that, in both scenarios, there exist succinct certificates ($O(n^2)$ in the size of the problem) and these certificates can be checked in subcubic (matrix multiplication) time. The certificates are based on grammar-based compression of paths (for positive instances) and on invariants represented as matrix constraints (for negative instances). Thus, CFL reachability lies in nondeterministic and co-nondeterministic *subcubic* time.

A natural question is whether faster algorithms for CFL reachability will lead to faster algorithms for combinatorial problems such as Boolean satisfiability (SAT). As a consequence of our certification results, we show that there cannot be a fine-grained reduction from SAT to CFL reachability for a conditional lower bound stronger than n^ω , unless the nondeterministic strong exponential time hypothesis (NSETH) fails.

Our results extend to related subcubic equivalent problems: pushdown reachability and two-way nondeterministic pushdown automata (2NPDA) language recognition. For example, we describe succinct certificates for pushdown non-reachability (inductive invariants) and observe that they can be checked in matrix multiplication time. We also extract a new hardest 2NPDA language, capturing the “hard core” of all these problems.

1 Introduction

Context-free reachability is a fundamental problem in interprocedural program analysis, verification of recursive programs, and database theory [20, 56, 38, 43, 8]. For a fixed context-free language (CFL) \mathcal{L} over an alphabet Σ , given a directed graph $G = (V, E)$, an edge-labeling function $\lambda : E \rightarrow \Sigma$, and two vertices $s, t \in V$, the \mathcal{L} -reachability problem asks if there is a path from s to t in G such that the word formed by concatenating the labels along the path belongs to \mathcal{L} . It is well-known that the problem can be solved in time cubic in the size of the graph for any fixed CFL. However, despite many years of efforts, we only know speedups by logarithmic factors (i.e., to $O(n^3/\log n)$) [45, 16], leading to a conjecture that no better algorithms are possible for this and several related problems [29]. In recent years, a number of results in fine-grained complexity give credence to the conjecture by demonstrating various conditional lower bounds for the problem [14, 1, 37], but even so the possibility of algorithms with running time n^ω or above has not been ruled out. Here, $\omega < 2.4$ is the matrix multiplication exponent [17, 54].

In this paper, we study the problem of *certifying* an instance of CFL reachability. Intuitively, this problem asks for easily verifiable proofs of inclusion or non-inclusion. Given a (positive or negative) instance of CFL reachability, we ask if there is an efficiently checkable proof that will convince anyone that the instance is indeed positive or negative.

Formally, a *certificate system* for CFL reachability consists of two algorithms (the checkers), one for positive instances and one for negative instance. Each checker takes as input an instance of the problem and an additional string (called the certificate) and accepts or rejects. The positive (resp. negative) checker is *complete* if for each positive (resp. negative) instance, there is a certificate that makes it accept, and *sound* if for each negative (resp. positive) instance, there is no certificate that makes it accept. Of course, since the instance can be decided in cubic time, a certificate system is non-trivial only if the checkers run in subcubic time (in the size of the instance).

Our main result shows the existence of **subcubic certificate systems for CFL reachability**: every positive or negative instance has a *quadratic* certificate and a checker that runs in $O(n^\omega)$ time.

- For a positive instance of the problem, a naive certificate is a path from s to t witnessing inclusion. Unfortunately, this is not an efficient certificate, since it is known that the shortest path can be exponentially long in the size of the graph. We show that the shortest path is well-compressible by a context free grammar of size $O(n^2)$ in the number of vertices of the graph. Moreover, given such a compressed representation, there is a checker verifying in time $O(n^2)$ that the grammar indeed encodes a witness path.
- For a negative instance of the problem, a certificate is an inductive invariant that demonstrates *non-reachability*. We show that such an inductive invariant can be represented as relations between a constant number of $n \times n$ matrices, and there is a checker verifying in time $O(n^\omega)$ that such an encoding does represent an inductive invariant. Additionally, if we allow randomization, there is a randomized checker running in $O(n^2)$ time.

Summing up, CFL reachability can be certified in subcubic time. In retrospect, the certificate system is simple but illuminates a conceptually new aspect of an old problem. Certificate systems make it possible to separate two possibly independent phases of computation, *finding* a solution to a computational problem and *verifying* it.

We consider complexity-theoretic implications. Impagliazzo and Paturi [31] introduced the strong exponential time hypothesis (SETH), which informally states that SAT has no algorithms better than exhaustive search. Over the years, SETH has become a fundamental assumption relative to which many fine-grained complexity results are proved [55]. For example, SETH implies current (quadratic) algorithms for orthogonal vectors or edit distance problems are optimal. A natural question is if SETH also implies that cubic algorithms for CFL reachability are optimal.

Our result shows that such a reduction would be very difficult to find. Carosino et al. [13] extended SETH to the nondeterministic strong exponential time hypothesis (NSETH), which states that there is no algorithm for Boolean tautology better than exhaustive search, even with nondeterministic guessing. They show that both proving and refuting NSETH imply breakthroughs in computational complexity. Our subcubic certification result implies that any conditional lower bound for CFL reachability from SAT and SETH will show that NSETH does not hold.

A model checking problem closely related to CFL reachability is pushdown reachability [8, 23]. Our results lead to a **subcubic certificate system for pushdown reachability** too, by extracting quadratic certificates from the standard saturation-based algorithm and the triplet construction for PDA to CFG conversion. Indeed, by exploiting fine-grained reductions between CFL reachability, pushdown reachability, the emptiness problem for pushdown automata, and the recognition problem for two-way nondeterministic pushdown automata (2NPDA), we show all these problems (as well as other related problems known in the literature) have subcubic certificate systems. Our constructions and reductions have several implications. First, succinct certificates for pushdown (non-)reachability checkable in subcubic time is a new observation; it can have potentially practical application in *checking* proofs of programs [40] and in “exports” of model checking such as certificate set analysis in trust management systems [32]. Second, our reductions lead to a new insight beyond certification. We identify a **new hardest 2NPDA language**, that is, a fixed 2NPDA language L_0 such that for every 2NPDA language L there is a homomorphism h such that $w \in L$ iff $h(w) \in L_0$. A different hardest language was previously found by Rytter [44] using language-theoretic techniques. However, our proof and reductions strengthen the link between 2NPDA language recognition and CFL reachability, pointing to the hardest instances of the latter.

Related work. In a quest to classify the complexity of problems in P, *fine-grained reductions* interlink the asymptotic running time of algorithms for various problems. A fine-grained reduction shows that a faster algorithm for one problem automatically implies a faster algorithm for another problem. Conversely, the existence of fine-grained reductions can be interpreted as *conditional lower bounds*: no faster algorithm exists, unless a state-of-the-art algorithm for a well-known problem is actually suboptimal. For example, a truly sub-quadratic algorithm for Orthogonal Vectors will lead to a $2^{(1-\varepsilon)n}$ -time algorithm for SAT, breaking SETH [55]. Similarly, the k -Clique conjecture states that no (randomized or deterministic) algorithm can detect a k -Clique on an n -vertex graph in time $O(n^{\frac{\omega k}{3}-\varepsilon})$ for $\varepsilon > 0$. Abboud et al. [1] show a reduction from the k -Clique problem to CFL *recognition*, giving a conditional lower bound of order n^ω and matching Valiant’s $\tilde{O}(n^\omega)$ upper bound for the problem [50]. This lower bound applies to CFL reachability as well. Chatterjee et al. [14], using Lee’s result [35], reduce Boolean matrix multiplication to Dyck- k reachability (for growing k), showing that faster algorithms for the latter avoiding matrix multiplication would be a breakthrough. Chatterjee and

Osang [15] show a similar reduction to PDA emptiness.

More broadly, a range of problems in formal languages are now being approached with tools from modern algorithms and complexity [21]. Our work contributes to this ongoing effort. Backurs and Indyk [4] and Bringmann, Grønlund, and Larsen [11] find SAT-based conditional lower bounds for regular expression matching problems. Oliveira and Wehar [18] show reductions between triangle finding, 3SUM, and the non-emptiness of intersection of two or three DFA. Potechin and Shallit [42] show a reduction from Orthogonal Vectors to the acceptance problem for (a subclass of) NFA and a reduction from triangle finding to (unary) NFA acceptance. Fernau and Krebs [22] establish conditional lower bounds for a variety of automata-theoretic problems beyond P. Wehar and co-authors have shown that faster algorithms for various intersection non-emptiness problems have consequences for structural complexity classes [53, 48, 19]. We discuss further related work in Section 6.

2 Context-free reachability and Dyck-2 reachability

Let \mathcal{L} be a fixed language. Given a directed graph $G = (V, E)$, an edge-labeling function $\lambda: E \rightarrow \Sigma$, and two vertices $s, t \in V$, the \mathcal{L} -reachability problem asks if there is a path from s to t (possibly repeating vertices and edges) such that the word formed by concatenating the labels along the path belongs to \mathcal{L} [56]. When \mathcal{L} is a fixed context-free language, the problem is called *CFL reachability*. CFL reachability plays a foundational role in several areas within computer science. To the best of our knowledge, it first appears in the work by Dolev, Even, and Karp [20] as the combinatorial core in the security analysis of a cryptographic protocol. Yannakakis [56] and Melski and Reps [38] elucidate the role of this problem in the context of database theory and interprocedural program analysis, respectively, providing in particular a historical sketch.

In formal language theory, \mathcal{L} -reachability and CFL reachability can be seen as providing an algorithmic perspective on the classic definition of rational index of a language [6, 41]. These problems have also been studied under the name “regular realizability” (see, e.g., [51, 52]).

For CFL reachability, without loss of generality, the fixed language can be assumed to be the *Dyck-2 language*. This is the language of balanced parentheses with two kinds of parenthesis symbols. Formally, it is the context free language over the alphabet $\{ (,), [,] \}$ defined by the following context-free grammar:

$$S \rightarrow SS \mid (S) \mid [S] \mid \varepsilon$$

The *Dyck-2 reachability* problem, denoted $D_2\text{Reach}$, is the \mathcal{L} -reachability problem when \mathcal{L} is the Dyck-2 language.

Claim 1. *Let (G, λ, s, t) be an instance of the CFL reachability problem. There is a linear-time reduction (in the bit-size of the input) to an instance (G', λ', s', t') of the Dyck-2 reachability problem.*

We call an algorithm *truly subcubic* if it has (worst-case) running time $O(n^{3-\varepsilon})$ for some constant $\varepsilon > 0$, where n denotes the bit length of the input. Practical implementations use a summarization-based $O(|V|^3)$ algorithm [43]; note that $|V| \leq n$. Using Rytter’s trick [45], Chaudhuri [16] shows that the \mathcal{L} -reachability problem is $O(|V|^3 / \log |V|)$ for any fixed context-free language. However, no truly subcubic algorithm is known for

this problem. The best known conditional lower bound for the problem is has order $|V|^\omega$. On the other hand, Dyck-1 reachability (the language of balanced parentheses with one kind of parentheses) can be solved in time $\tilde{O}(|V|^\omega)$ [9, 10, 37], matching best conditional lower bounds.

3 Certificates for reachability and non-reachability

In this section we show that, while truly subcubic algorithms for Dyck-2 reachability are not known, solutions to Dyck-2 reachability have small and efficiently checkable certificates.

An instance (G, λ, s, t) of $D_2\text{Reach}$ is a yes-instance if there is a walk from s to t labeled with a string from Dyck-2, and a no-instance otherwise.

Definition 2. We say that $D_2\text{Reach}$ has *subcubic certificates* for yes-instances (respectively, no-instances) if, for some real number $\varepsilon > 0$, there is an algorithm M and a function $p(x) = O(x^{3-\varepsilon})$ such that for every instance (G, λ, s, t) of $D_2\text{Reach}$:

(completeness) if the instance is a yes-instance (respectively, no-instance), then there is a string u of length $p(|V|)$, called a *certificate*, such that M accepts (G, λ, s, t, u) in $p(|V|)$ time, and

(soundness) if the instance is a no-instance (respectively, yes-instance), then for every string u of length $p(|V|)$, the algorithm M rejects (G, λ, s, t, u) in $p(|V|)$ time.

(Note that the running time of M is subcubic in $|V|$, which is at most the bit size of the instance, and *not* in the size of the certificate.) That is, a subcubic certificate for yes- and no-instances allows us to verify, given the additional certificate, whether an instance of $D_2\text{Reach}$ is a positive or a negative instance in sub-cubic time.

We will refer extensively to *walks* in labelled directed graphs. For a labelled directed graph $(V, E, \lambda: E \rightarrow \{(\cdot, \cdot), [\cdot, \cdot]\})$, a walk from $u \in V$ to $v \in V$ is a sequence of edges $\pi := e_0 \dots e_k$ from E , for $k \geq 0$, such that for each $i \in \{1, \dots, k\}$, edge e_{i-1} arrives at the same vertex that edge e_i departs from, and moreover e_0 departs from u and e_k arrives at v . This walk is *valid* if the word $\lambda(e_0) \dots \lambda(e_k)$ belongs to the Dyck-2 language. A subwalk of a walk $e_0 \dots e_k$ is a contiguous subsequence $e_i \dots e_j$ of edges, possibly empty.

3.1 Certificates for yes-instances: compressed walks

We describe our certificate system for yes-instances of $D_2\text{Reach}$. These certificates are witnesses for reachability. We fix an instance of $D_2\text{Reach}$: $G = (V, E)$ a directed graph, $\lambda: E \rightarrow \{(\cdot, \cdot), [\cdot, \cdot]\}$ an edge-labeling function, and $s, t \in V$ source and target vertices.

A first attempt is to provide a valid walk as a certificate (witness). However, it is well-known that the shortest valid walk can be exponential in the size of the input, namely it can be of length $\exp \Theta(|V|^2 / \log |V|)$, and this bound is tight [41]. (For an intuition, one can think of a pushdown automaton accepting only words of length exponential in its size and longer.) The main observation to get subcubic certificates is that there is always some valid walk (including the shortest one in particular) that is well-compressible and that has a small representation ($O(|V|^2)$ in the size of the graph) and it is efficient to check (in time $O(|V|^2)$) that such a compressed walk is indeed a valid walk. Moreover, for every no-instance, one cannot get any valid walks, compressed or otherwise.

The following definition “inlines” the concept of a straight-line program, which is an “acyclic” context-free grammar that generates one word only. Straight-line programs are at the core of general-purpose compression algorithms such as LZ77 (see, e.g., [36]).

Definition 3. For an instance of $D_2\text{Reach}$, denote by \vec{V}^2 a fresh copy of the set V^2 , written as $\vec{V}^2 = \{\vec{uv} \mid (u, v) \in V^2\}$. A *walk scheme* is a context-free grammar with the set of terminal symbols E , a set of nonterminal symbols $\text{NT} \subseteq \vec{V}^2$, and the axiom $\vec{st} \in \text{NT}$, where:

- for each nonterminal $\vec{uv} \in \text{NT}$ there is exactly one production, which moreover has the form:
 - (a) $\vec{uv} \rightarrow \vec{uw} \vec{wv}$ for some $w \in V$, or
 - (b) $\vec{uv} \rightarrow e \vec{xy} f$ for some edges $e = (u, x) \in E$ and $f = (y, v) \in E$ with $\lambda(e) \cdot \lambda(f) \in \{(), []\}$, or
 - (c) $\vec{uu} \rightarrow \varepsilon$ for some $u \in V$, and
- the directed graph with vertices NT and the following set of edges is acyclic:

$$\left\{ (\vec{ab}, \vec{cd}) \mid \vec{cd} \text{ occurs on the right-hand side of the production of } \vec{ab} \right\}. \quad (1)$$

Proposition 4. *Every walk scheme has size $O(|V|^2)$ and bit size $O(|V|^2 \log |V|)$.*

Theorem 5. *The following statements hold:*

- *An instance of $D_2\text{Reach}$ is a yes-instance if and only if there exists a walk scheme for it.*
- *There is a deterministic algorithm that runs in time $O(|V|^2)$ and decides if a given grammar is a walk scheme for a given instance of $D_2\text{Reach}$.*

For the proof of Theorem 5, we need the following auxiliary result.

Lemma 6. *Let \mathcal{G} be a context-free grammar with $L(\mathcal{G}) \neq \emptyset$. Suppose \mathcal{G} contains more than one production with the same nonterminal on the left-hand side. Then by removing all of them but one we can obtain a grammar \mathcal{G}' with $L(\mathcal{G}') \neq \emptyset$.*

Proof of Theorem 5. We split the proof into three parts.

Soundness. We first suppose that for a given instance of $D_2\text{Reach}$ there exists a walk scheme, \mathcal{W} , and show that the instance must be a yes-instance. Consider the directed graph from the acyclicity condition in the definition of walk schemes, denote it D . We will consider all vertices of D , i.e., nonterminals from NT , in any reversed topological ordering. In other words, whenever \vec{cd} occurs on the right-hand side of the production of \vec{ab} , we will consider \vec{cd} before \vec{ab} . We will show by induction that, for every $\vec{uv} \in \text{NT}$, the (one) word generated by \vec{uv} is a valid walk from u to v . (Recall that a walk is valid if it is labelled by a Dyck-2 word.) Indeed, it suffices to consider the three types of productions:

- (a) for a production of the form $\overrightarrow{uv} \rightarrow \overrightarrow{uw} \overrightarrow{wv}$, we know from the inductive hypothesis that \overrightarrow{uw} generates a valid walk from u to w , and \overrightarrow{wv} a valid walk from w to v , so their concatenation is a valid walk from u to v ;
- (b) for a production of the form $\overrightarrow{uv} \rightarrow e \overrightarrow{xy} f$ with edges $e = (u, x) \in E$ and $f = (y, v) \in E$, we know from the inductive hypothesis that \overrightarrow{xy} generates a valid walk from x to y , and since $\lambda(e) \cdot \lambda(f) \in \{(), \square\}$, the result of the concatenation is a valid walk from u to v ;
- (c) finally, productions of the form $\overrightarrow{uu} \rightarrow \varepsilon$ correspond to trivial valid walks (containing no edges) and represent the induction base.

As the axiom of the grammar \mathcal{W} is \overrightarrow{st} , we conclude that there is a valid walk from s to t , which means that the instance of D_2Reach we consider is a yes-instance.

Completeness. In the converse direction, let us prove that every yes-instance of D_2Reach has a walk scheme. Consider such an instance, (G, λ, s, t) , and consider a walk from s to t , call it π . We construct a walk scheme in several steps.

First consider a context-free grammar \mathcal{G} with the set of terminal symbols E , set of nonterminal symbols $\overrightarrow{V^2}$, and axiom \overrightarrow{st} . The set of productions is determined as follows. For each nonterminal $\overrightarrow{uv} \in \overrightarrow{V^2}$, we include all productions of the form:

- $\overrightarrow{uv} \rightarrow \overrightarrow{uw} \overrightarrow{wv}$ for all $w \in V$;
- $\overrightarrow{uv} \rightarrow e \overrightarrow{xy} f$ where $e = (u, x) \in E$ and $f = (y, v) \in E$ such that $\lambda(e) \cdot \lambda(f) \in \{(), \square\}$;
- $\overrightarrow{uu} \rightarrow \varepsilon$ for all $u \in V$.

Induction on the structure of π shows that $\pi \in L(\mathcal{G})$, so $L(\mathcal{G}) \neq \emptyset$.

We can now prune the set of productions of the grammar \mathcal{G} using Lemma 6, as well as apply standard procedures of removing useless (non-productive or unreachable) nonterminals in context-free grammars (see, e.g., [30, Section 7.1]). We perform these steps until all three have no effect on the grammar. The resulting grammar \mathcal{W} satisfies all conditions in the definition of walk schemes, except possibly the acyclicity condition. We claim that \mathcal{W} must satisfy that condition too. Indeed, the transformations applied so far ensure that $L(\mathcal{W}) \neq \emptyset$. Let $\text{NT} \subseteq \overrightarrow{V^2}$ be the set of nonterminals of \mathcal{W} . Assume for the sake of contradiction that the directed graph with vertices NT and edges (1) contains a directed cycle. Let $\overrightarrow{ab} \in \text{NT}$ be a vertex on this cycle. Since all nonterminals of \mathcal{W} are reachable and productive, there exists a valid parse tree with respect to \mathcal{W} that contains a node labelled by \overrightarrow{ab} . By definition of the graph, and since every nonterminal in \mathcal{W} has exactly one production, this node has a descendant labelled with \overrightarrow{ab} . By the same reasoning, this descendant also has a descendant labelled with \overrightarrow{ab} , etc., which cannot be the case as the tree is finite. This contradiction means that the graph must be acyclic, so \mathcal{W} is in fact a walk scheme.

Verification algorithm. The condition $\text{NT} \subseteq \overrightarrow{V^2}$ and the choice of the axiom can be checked in time $O(|V|^2)$. The fact that there is exactly one production per nonterminal can be checked under the same time constraints; and so can the form of these productions and compatibility with the instance of D_2Reach . Finally, depth-first search-based topological sort procedure can be used to detect the existence of directed cycles; it runs in time linear in the number of edges, which is at most $|V|^2$. \square

Remark 7. There is nothing special about Dyck-2 in the construction, and a similar certificate can be constructed for any fixed CFG.

We already mentioned a link to compressed words above. Our proof of Theorem 5 finds a context-free grammar that generates exactly one word and has $O(|V|^2)$ nonterminals in Chomsky normal form. Importantly, while it is in general a PSPACE-complete problem to decide whether such a compressed word is accepted by a pushdown automaton (see, e.g., the survey [36, section 9.4] and references therein), our grammar has special structure, leading to an efficient verification algorithm.

3.2 Certificates for no-instances: inductive invariants

Fix an instance of D_2Reach . For ease of notation, we will assume that $V = \{1, \dots, |V|\}$. A certificate for no-instances will be a *separator*, as defined next. Such a certificate is essentially an inductive invariant, certifying non-reachability.

Let $A_{\langle}, A_{\lceil}, A_{\rangle}, A_{\rfloor}$ be four 0–1 matrices of size $|V| \times |V|$ that are adjacency matrices for the graph G restricted to sets of edges with labels $\langle, \lceil, \rangle, \rfloor$, respectively.

For a nonnegative integer matrix N , denote by $\text{bool}(N)$ the matrix obtained from N by replacing every nonzero element by 1. Let I denote the $|V| \times |V|$ identity matrix. We write $A \leq B$ for matrices $A = (a_{ij})$ and $B = (b_{ij})$ of the same size whenever $a_{ij} \leq b_{ij}$ for all i, j .

Definition 8. A *separator* for an instance of D_2Reach is a sextuple of $|V| \times |V|$ matrices, $(M_S, M_{SS}, M_{\langle S}, M_{\lceil S}, M_{\rangle S}, M_{\rfloor S})$, where all entries belong to $\{0, 1, \dots, |V|^2\}$, and moreover all entries of M_S belong to $\{0, 1\}$, and such that the following ten conditions are satisfied:

$$\begin{aligned} I &\leq M_S, & A_{\langle} \cdot M_S &= M_{\langle S}, & A_{\lceil} \cdot M_S &= M_{\lceil S}, \\ M_S \cdot M_S &= M_{SS}, & M_{\langle S} \cdot A_{\rangle} &= M_{\langle S}, & M_{\lceil S} \cdot A_{\rfloor} &= M_{\lceil S}, \\ \text{bool}(M_{SS}) &\leq M_S, & \text{bool}(M_{\langle S}) &\leq M_S, & \text{bool}(M_{\lceil S}) &\leq M_S, & \text{and } (M_S)_{s,t} &= 0, \end{aligned} \tag{2}$$

where s and t are the source and target vertex in the instance of D_2Reach .

Proposition 9. *Every separator has $O(|V|^2)$ entries and bit size $O(|V|^2 \log |V|)$.*

Theorem 10. *The following statements hold:*

- *An instance of D_2Reach is a no-instance if and only if there exists a separator for it.*
- *There is a deterministic algorithm that runs in time $O(|V|^\omega)$ and decides if a given sextuple of $|V| \times |V|$ matrices is a separator for a given instance of D_2Reach .*

- *There is a randomized algorithm that runs in time $O(|V|^2)$ and decides if a given sextuple of $|V| \times |V|$ matrices is a separator for a given instance of $D_2\text{Reach}$. In the case it is, the algorithm never errs; otherwise the algorithm flags an issue with probability ≥ 0.5 .*

Proof. We split the proof into four parts.

Completeness. First consider a no-instance of $D_2\text{Reach}$. Take the matrix $M_S = (m_{ij})$, where each m_{ij} is 1 if there is a valid walk from vertex i to vertex j . It is clear that $m_{st} = 0$, because the instance is a no-instance. We now show that picking the other matrices $M_{SS}, M_{(S)}, M_{[S]}, M_{(S)}, M_{[S]}$ so that all the five matrix equalities among the constraints (2) are satisfied leads to the satisfaction of the remaining (four) inequality constraints. Indeed:

- $I \leq M_S$ because for each vertex i the empty walk from i to i is valid;
- $\text{bool}(M_{SS}) \leq M_S$ because the concatenation of two valid walks is a valid walk;
- $\text{bool}(M_{(S)}) \leq M_S$ and $\text{bool}(M_{[S]}) \leq M_S$ because every walk $e \cdot \pi \cdot e'$ is valid whenever π is valid and e and e' are labelled by a matching pair of parentheses, either $(,)$ or $[,]$.

This shows that there is a separator for each no-instance.

Soundness. In the converse direction, consider an arbitrary instance of $D_2\text{Reach}$. We show that for every valid walk π from a vertex u to a vertex v in the graph, all separators must satisfy the condition $m_{uv} = 1$ where $M_S = (m_{ij})$. (It then follows that yes-instances have no separators.) We use induction on the label of walk π , which is simply the concatenation of individual edge labels:

- The base case is the empty label, ε . The walk π must then be the empty walk, from some vertex u to itself. We recall that $I \leq M_S$ for every separator; so indeed m_{ii} must be set to 1 for all vertices i , and for the chosen vertex $i = u$ in particular.
- If the walk π is labelled by $\alpha \cdot \beta$, where both α and β are nonempty Dyck-2 words, then there exists a vertex w such that $\pi = \pi' \cdot \pi''$ and π' and π'' are valid walks from u to w and from w to v , respectively. By the inductive hypothesis, $m_{u,w} = m_{w,v} = 1$. Since $\text{bool}(M_{SS}) = \text{bool}(M_S \cdot M_S) \leq M_S$, we conclude that $m_{u,v} = 1$ in this case as well.
- Finally, suppose the label of the walk π is (α) , for some Dyck-2 word α . (The case $[\alpha]$ is analogous.) Then $\pi = e \cdot \pi' \cdot f$, where e and f are individual edges, say from u to u' and from v' to v (for some $u', v' \in V$), and π' is a valid walk from u' to v' . The edges $e = (u, u')$ and $f = (v', v)$ have labels $($ and $)$, respectively. By the inductive hypothesis, $m_{u'v'} = 1$. We now observe that $\text{bool}(M_{(S)}) = \text{bool}(M_{(S)} \cdot A_{\text{)}} = \text{bool}(A_{(} \cdot M_S \cdot A_{\text{)}}) \leq M_S$. On the left-hand side, the matrix product has a positive entry in position uv , because $(A_{(})_{u,u'} = (A_{\text{)}}_{v',v} = 1$ by the definition of $A_{(}$ and $A_{\text{}}$. Therefore $m_{uv} = 1$.

This concludes the proof of the first assertion of the theorem.

Deterministic algorithm. The algorithm from the second assertion of the theorem verifies all conditions in the definition of separator directly. This means in particular five matrix multiplications where the factors are matrices with elements from $\{0, \dots, |V|\}$ (worst-case time $O(|V|^\omega)$), four inequalities between individual matrices (worst-case time $O(|V|^2)$), and a single equality constraint on one of the entries (constant time).

Remark. This algorithm reduces the verification of separators to 5 matrix multiplications over the nonnegative integers. While this result has complexity-theoretic consequences (see Section 4 below), it may appear unsatisfactory, as many theoretical algorithms for fast matrix multiplication are impractical. This brings the randomized algorithm to the fore.

Randomized algorithm. The algorithm from the final assertion of the theorem is the same as the previous one, except that instead of *computing* matrix multiplication it runs Freivalds' algorithm for *verifying* matrix multiplication [24].

Recall that Freivalds' algorithm for verifying $A \cdot B = C$ for some $n \times n$ matrices A , B , and C proceeds by picking a 0–1 vector $u \in \{0, 1\}^n$ uniformly at random and checking if $A \cdot (Bu) = Cu$. The algorithm runs in $O(n^2)$ time and has error probability $1/2$. The properties of the algorithm are transferred directly to give a $O(|V|^2)$ bound. Since we have five products to check, we reduce the error probability in an individual check to $1/16$ by running it 4 times, so that the overall error probability is at most $5/16 \leq 1/2$. \square

Remark 11. For the deterministic verification algorithm, it suffices to specify the 0–1 matrix M_S only, because the other five matrices can be computed in time $O(|V|^\omega)$ from it.

Remark 12. Once again, there is nothing special about the Dyck-2 language in our certificate system. One can readily see that the conditions we impose on separators correspond to the following context-free grammar for the Dyck-2 language:

$$S \rightarrow SS \mid P \mid Q \mid \varepsilon \qquad P \rightarrow (S \qquad Q \rightarrow [S .$$

Replacing this grammar with a different one, we obtain a certificate system (for no-instances) for the CFL reachability problem where the fixed CFL is represented by any fixed CFG.

Remark 13. In a model of computation with unit-cost integer arithmetic, integer matrix multiplication can be verified in deterministic time $O(n^2)$ [33]. For RAM with $O(\log n)$ -bit arithmetic operations, derandomization of Freivalds' algorithm is an open problem even in the nondeterministic setting. However, if the number of errors in the product is guaranteed to be $O(n^{2-\varepsilon})$, then a deterministic $O(n^{3-\varepsilon})$ -time algorithm is known [34].

4 Complexity implications

Complexity-theoretic summary of Section 3. Leaving out sharper bounds on certificate size and $\text{polylog}(n)$ factors (required in the Turing model), Theorems 5 and 10 imply:

Theorem 14. $\text{D}_2\text{Reach} \in \text{NTIME}(n^2) \cap \text{coNTIME}(n^\omega) \cap \text{co-MATIME}_1(n^2)$.

For this summary, we recall (cf. [49]) that $L \in \text{MATIME}_1(t)$ (Merlin-Arthur time, introduced by Babai [3]) iff there exists a deterministic machine M that takes inputs x, y, z where $|y| = |z| = O(t(|x|))$, runs in time $O(t(|x|))$, and such that for every x ,

$$x \in L \Rightarrow \exists y. \Pr_z[M(x, y, z) \text{ accepts}] = 1, \quad x \notin L \Rightarrow \forall y. \Pr_z[M(x, y, z) \text{ accepts}] \leq 1/2,$$

where the probability is with respect to the uniform distribution of z in $\{0, 1\}^{t(|x|)}$. Finally, $\text{co-MATIME}_1(t)$ is the class of *complements* of languages in $\text{MATIME}_1(t)$.

Fine-grained complexity of D_2Reach . Fine-grained complexity research shows that even small improvements in (the exponent of) the running time of many algorithmic problems, such as orthogonal vectors or edit distance, would automatically give faster algorithms for Boolean satisfiability, SAT [55]. Would improvements over Chaudhuri's $O(n^3/\log n)$ -time algorithm for D_2Reach also have consequences for SAT ? Here we show that subcubic certificates give an answer to this question.

In fine-grained complexity, perhaps the most influential hypothesis, and the ultimate source of many lower bounds, is the *strong exponential-time hypothesis* (SETH) [31], stating (roughly) that there is no algorithm for SAT better than exhaustive enumeration. The *non-deterministic strong exponential-time hypothesis* (NSETH) [13] extends it further.

Hypothesis 15 (SETH). *For every $\varepsilon > 0$, there exists a k so that $k\text{-SAT}$ is not in $\text{DTIME}[2^{n(1-\varepsilon)}]$, where $k\text{-SAT}$ is the language of all satisfiable Boolean formulas in $k\text{-CNF}$.*

Hypothesis 16 (NSETH). *For every $\varepsilon > 0$, there exists a k so that $k\text{-TAUT}$ is not in $\text{NTIME}[2^{n(1-\varepsilon)}]$, where $k\text{-TAUT}$ is the language of all Boolean tautologies in $k\text{-DNF}$.*

In both hypotheses, n is the number of variables. It is unknown whether SETH and NSETH are true. NSETH implies SETH , and SETH implies $\text{P} \neq \text{NP}$. Carmosino et al. [13] explore consequences of NSETH and show that both proving and refuting it would lead to interesting consequences. In particular, NSETH implies the *absence* of fine-grained reductions from SAT to a number of problems and $\neg\text{NSETH}$ implies circuit lower bounds.

It turns out that, because of our subcubic certificate systems (Section 3), there exists no fine-grained reduction from SAT (as well as from any SETH -hard problem) to D_2Reach that would imply hardness beyond n^ω , unless NSETH fails.

Because of space constraints, we relegate the formal definition of fine-grained reductions to Appendix D. Intuitively, a fine-grained reduction from $(L, t(n))$ to $(\text{D}_2\text{Reach}, n^c)$ means that, for every $\varepsilon > 0$, an $O(n^{c-\varepsilon})$ -time algorithm for D_2Reach implies a $O(t(n)^{1-\delta})$ algorithm for problem L for some $\delta = \delta(\varepsilon) > 0$. This is not unlike usual Turing reductions (allowing multiple queries), tracking the precise exponents in the running time bounds. The following result is a consequence of Theorem 14.

Theorem 17. *Unless NSETH fails, there is no fine-grained reduction from $(\text{SAT}, 2^n)$ to $(\text{D}_2\text{Reach}, n^{\omega+\gamma})$ for any $\gamma > 0$.*

5 Certificates for pushdown non-reachability

While CFL reachability is a central problem in program analysis, an analogous problem in model checking is *pushdown reachability* [8, 23, 7, 47], formalized as follows.

We are given a pushdown automaton (PDA) $\mathcal{P} = (Q, \Gamma, \Delta)$, where Q is a finite set of states, Γ is a finite alphabet of stack symbols, and $\Delta \subseteq (Q \times \Gamma) \times (Q \times \Gamma^{\leq 2})$ is a set of transitions, and an initial configuration $(q_0, \gamma_0) \in Q \times \Gamma$. We are additionally given a regular set of configurations R specified by a \mathcal{P} -automaton: this is a usual, ε -free nondeterministic finite automaton (NFA) over the alphabet Γ in which the set of control states is $S \supseteq Q$ and the transition relation is $\delta \subseteq S \times \Gamma \times S$. A set of final states, $F \subseteq S$, is usually taken to be disjoint from Q . Such a \mathcal{P} -automaton is said to accept a configuration $(q, w) \in Q \times \Gamma^*$ of the PDA \mathcal{P} iff there is a walk from control state q to some $\bar{q} \in F$ labelled by the word w ; in other words, if w is accepted by this NFA when started from q as initial state. We ask if the PDA \mathcal{P} has a run from (q_0, γ_0) to some configuration from R .

We adapt our certificate system to pushdown reachability. For yes-certificates of size $O(|\Gamma||S|^2)$, we can convert the PDA to an equivalent CFG using the standard triplet construction (see, e.g., [30, Chapter 6]) and repeat the second half of the completeness argument from Subsection 3.1. Explicitly, a certificate is a “sub-grammar” of this CFG that is a straight-line program.

We now show how to certify that a given initial configuration *cannot reach* any configuration from a given regular set R . The classic saturation algorithm for computing $\text{Pre}^*(R)$, the set of (reflexive, transitive) predecessors of configurations in R , takes a \mathcal{P} -automaton \mathcal{A} as input and iteratively adds transitions to it by the following rule:

$$\mathcal{P} \text{ has transition } (p, A) \rightarrow (q, w), \mathcal{A} \text{ has walk } q \xrightarrow{w} s \Rightarrow \text{add transition } p \xrightarrow{A} s \text{ to } \mathcal{A}. \quad (3)$$

By the following claim, saturation under (3) implies overapproximation of $\text{Pre}^*(R)$. The converse inclusion is more subtle and will not be required.

Claim 18 (see, e.g., Carayol and Hague [12, Section 3.2]). *A \mathcal{P} -automaton \mathcal{A} accepts all configurations from $\text{Pre}^*(R)$ if (i) it contains all transitions of the original \mathcal{P} -automaton and (ii) it is saturated, i.e., applying rule (3) does not change the transition relation.*

Our certificate system for non-reachability relies on the observation that the update rule (3) can be expressed using matrix multiplication. A *certificate* is a finite family of matrices, $M^A, M^{A,B}, M_1^{A,B,C}, M_2^{A,B,C}$, for all $A, B, C \in \Gamma$, satisfying the following conditions:

$$\begin{aligned} P^A &\leq M^A, & (M^{\gamma_0})_{q_0, f} &= 0 & \text{for all } f \in F, \\ T^{A, \varepsilon} &\leq M^A, \\ \text{bool}(M^{A,B}) &\leq M^A, & M^{A,B} &= T^{A,B} \cdot M^B, \\ \text{bool}(M_2^{A,B,C}) &\leq M^A, & M_1^{A,B,C} &= T^{A,BC} \cdot M^B, & M_2^{A,B,C} &= M_1^{A,B,C} \cdot M^C, \end{aligned} \quad (4)$$

where we assume with no loss of generality that $S = \{1, \dots, |S|\}$ and denote by P^A the A -transition matrix of the original \mathcal{P} -automaton and, for all $A \in \Gamma, w \in \Gamma^{\leq 2}$, by $T^{A,w} = (t_{ij}^{(A,w)})$ the 0–1 matrix of size $|S| \times |S|$ in which $t_{ij}^{(A,w)} = 1$ if $i, j \in Q$ and \mathcal{P} contains a transition $(i, A) \rightarrow (j, w)$. The following proposition summarises the properties of this system:

Proposition 19. *Certificates have $O(|\Gamma|^3|S|^2)$ entries. An instance of PDA emptiness is a no-instance iff there exists a certificate for it. The conditions can be verified by a deterministic algorithm with running time $O(|\Gamma|^3|S|^\omega)$ or a randomized algorithm with running time $O(|\Gamma|^3|S|^2)$ that accepts valid certificates with probability one and rejects invalid ones with probability ≥ 0.5 .*

Matrix constraints of Eq. (4) define a *backwards invariant* for the pushdown system \mathcal{P} in question, an overapproximation of the set of configurations from which R is reachable.

6 Discussion: Fine-grained landscape and a hardest $D_2\text{Reach}$ instance

In interprocedural program analysis, the lack of algorithms with running time $O(n^{3-\varepsilon})$ is referred to as “the cubic bottleneck”. Heintze and McAllester [29] captured this phenomenon by the class of “2NPDA-complete” problems. Here “2NPDA” stands for *two-way nondeterministic pushdown automata*, a model of computation that extends standard PDA with the ability to move back and forth on the (read-only) input tape [2]. A problem is 2NPDA-complete (following Neal [39]) if it is subcubic equivalent to *2NPDA recognition*: given a word, does it belong to the language of a fixed 2NPDA. Heintze and McAllester show a number of 2NPDA-complete problems, including ground monadic rewriting reachability (see also [39]), data flow reachability, control flow reachability, and certain (non-)typability problems. Melski and Reps [38] show a reduction from CFL reachability to data flow reachability and set constraints (and thus to 2NPDA recognition) and a reverse reduction from data flow reachability to an instance of CFL reachability where the language is *not* fixed.

The following result appears to be folklore but is not found in the literature, strengthening the reduction of Melski and Reps to show hardness of CFL reachability for the *fixed* Dyck-2 language. The equivalence between problems (1) and (2) is sketched by Chaudhuri [16]. While we state the result for PDA emptiness, one can equivalently (or additionally) state it for pushdown reachability. We provide full proofs in the appendix.

Proposition 20. *The following problems either all have truly subcubic algorithms, or none of them do: (1) 2NPDA language recognition, (2) PDA language emptiness, and (3) $D_2\text{Reach}$.*

Proof (sketch). We show three reductions:

- In 2NPDA recognition to PDA emptiness, each control state of the PDA remembers the position of the 2NPDA on the input tape and the control state of the 2NPDA. The size of PDA is linear in the length of the input word, because the 2NPDA is fixed.
- In PDA emptiness to $D_2\text{Reach}$, the graph mimics the transition diagram of the PDA. Stack symbols from Γ are encoded by sequences of opening parentheses of two kinds of length $\lceil \log |\Gamma| \rceil$. Push transitions are modelled by sequences of edges with these labels, and pop transitions by sequences with matching closing parentheses. The reduction is linear-time, because the bit size of the PDA accounts for the $\log |\Gamma|$ factor.
- In the last reduction, we give a fixed 2NPDA that solves $D_2\text{Reach}$. The 2NPDA guesses a path through the graph, maintaining at the bottom of the stack a sequence $\sigma \in \{(\, [\}^*$, and the current vertex at the top of the stack. The length of the input word is proportional to the bit size of the graph (adjacency lists). \square

As a corollary, all of these problems have subcubic certificate schemes, and an analogue of Theorem 14 holds for them too (worked out for PDA emptiness in Section 5). Theorem 17 on the absence of SETH-hardness also extends to PDA emptiness and 2NPDA recognition.

For upper bounds, note that 2NPDA recognition is solvable in time $O(|w|^3/\log |w|)$ [45], and language emptiness for PDA in time $O(n^3/\log n)$ ¹.

We observe that the hardness of 2NPDA recognition is witnessed by a single “hardest” 2NPDA language: recognition for an arbitrary 2NPDA can be reduced to a single 2NPDA. Suppose some 2NPDA \mathcal{A} over Σ is given and the input to 2NPDA recognition for \mathcal{A} is a word w . Applying our cycle of reductions from Proposition 20 (to PDA emptiness, then to CFL reachability, and then back to 2NPDA recognition), we get another word $u = u(\mathcal{A}, w)$ and a 2NPDA $\mathcal{B} = \mathcal{B}(\mathcal{A}, w)$ such that \mathcal{B} accepts u iff \mathcal{A} accepts w . But \mathcal{B} in fact doesn’t depend on \mathcal{A} or w , because it is a fixed 2NPDA for D_2Reach . One refers to such languages as *hardest* 2NPDA languages, because the recognition problem for $L(\mathcal{B})$ cannot be easier than the recognition problem for any 2NPDA language L . The following theorem states this result in language-theoretic terms. (Recall that a homomorphism is a mapping, say $h: \Sigma^* \rightarrow \Sigma_0^*$, such that $h(uv) = h(u)h(v)$ for all $u, v \in \Sigma^*$.)

Theorem 21. *There exists a 2NPDA \mathcal{A}_0 over an input alphabet Σ_0 with the following property: for every 2NPDA \mathcal{A} over every finite Σ there is a homomorphism $h: \Sigma^* \rightarrow \Sigma_0^*$ such that, for all $w \in \Sigma^+$, $w \in L(\mathcal{A})$ if and only if $h(w) \in L(\mathcal{A}_0)$.*

Essentially, $\mathcal{B} = \mathcal{A}_0$. Working out the details shows that the mapping $u(\mathcal{A}, \cdot)$ can be made a homomorphism for every \mathcal{A} . This requires an appropriate encoding for inputs to \mathcal{A}_0 .

Remark 22. Rytter [44] showed there is a fixed hardest 2NPDA language L_0 ,² based on the classic hardest context-free language by Greibach [28]. Theorem 21 identifies a different hardest 2NPDA language. In contrast with Rytter’s proof, our construction is self-contained and does not depend on Greibach’s hardest CFL. Instead, our new hardest 2NPDA language is an encoding of a restricted version of Dyck-2 Reachability.

We now describe the hardest language $L(\mathcal{A}_0)$. The alphabet is $\Sigma_0 = \{ (,), [,], \#, 1, -, * \}$. The language contains only words of the form

$$\#\ell_1 o_1 * \ell_2 o_2 * \dots * \ell_q o_q \# \ell_{q+1} o_{q+1} \dots \# \dots \ell_m o_m \tag{5}$$

and the membership of such words in the language is determined as follows. Consider a directed graph $G = (V, E)$ with $V = \{1, \dots, n\}$ where n is the number of *blocks* separated by the vertex marker $\#$. An edge $e = (i, j)$ belongs to E if and only if the i th block has a subword $\ell_p o_p$ with $\ell_p \in \{ (,), [,] \}$, $o_p = 1^k$ or $o_p = -1^k$ where $j = i + k$ and this subword is preceded and followed by symbols from $\{ \#, * \}$ or tape endmarker. The edge label is in this case $\lambda(e) = \ell_i$. (If for some i and k the index j is “off the tape”, the tape endmarker counts as one virtual vertex and then the counting reverses the direction, “reflecting” off the endmarker.) The word belongs to $L(\mathcal{A}_0)$ if and only if $(G, \lambda, 1, n)$ is a yes-instance

¹The reduction of Proposition 20, combined with Chaudhuri’s algorithm for CFL reachability [16], implies a $O(n^3/\log n)$ bound for PDA emptiness where n is the bit size of the input. (We give a sketch in Appendix F.) In contrast, “textbook” algorithms for PDA emptiness go through equivalent context-free grammars [30], for which a cubic blow-up is unavoidable in the worst case [26].

²Actually, Rytter only proves that, for all $w \in \Sigma^+$, one has $w \in L$ iff $h(w\$) \in L_0$.

of $D_2\text{Reach}$, i.e., if G contains a walk from 1 to n labelled by a word from the Dyck-2 language.

To sum up, this restricted version of 2NPDA recognition is the “hard core” of the problem: by Theorems 21 and 20, in order to find subcubic algorithms for $D_2\text{Reach}$, it suffices to handle instances obtained from it (exploiting any structural properties). PDA emptiness and $D_2\text{Reach}$ are already hard for sparse graphs: a truly subcubic algorithm for either problem restricted to graphs with a linear number of edges would already result in a breakthrough algorithm for 2NPDA recognition.

Acknowledgements.

We thank Sayan Bhattacharya and Karl Bringmann for interesting discussions. Philipp Schepper is supported by the European Research Council (ERC) consolidator grant no. 725978 SYSTEMATICGRAPH. Rupak Majumdar was funded in part by the Deutsche Forschungsgemeinschaft project 389792660-TRR 248 and by the European Research Council under the Grant Agreement 610150 (ERC Synergy Grant ImPACT).

References

- [1] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. If the current clique algorithms are optimal, so is Valiant’s parser. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 98–117. IEEE Computer Society, 2015.
- [2] Alfred V. Aho, John E. Hopcroft, and Jeffrey D. Ullman. Time and tape complexity of pushdown automaton languages. *Information and Control*, 13(3):186–206, 1968.
- [3] László Babai. Trading group theory for randomness. In Robert Sedgewick, editor, *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6-8, 1985, Providence, Rhode Island, USA*, pages 421–429. ACM, 1985. URL: <https://doi.org/10.1145/22145.22192>, doi:10.1145/22145.22192.
- [4] Arturs Backurs and Piotr Indyk. Which regular expression patterns are hard to match? In Irit Dinur, editor, *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 457–466. IEEE Computer Society, 2016. URL: <https://doi.org/10.1109/FOCS.2016.56>, doi:10.1109/FOCS.2016.56.
- [5] Daniel Bienstock, Neil Robertson, Paul D. Seymour, and Robin Thomas. Quickly excluding a forest. *J. Comb. Theory, Ser. B*, 52(2):274–283, 1991. URL: [https://doi.org/10.1016/0095-8956\(91\)90068-U](https://doi.org/10.1016/0095-8956(91)90068-U), doi:10.1016/0095-8956(91)90068-U.
- [6] Luc Boasson, Bruno Courcelle, and Maurice Nivat. The rational index: A complexity measure for languages. *SIAM J. Comput.*, 10(2):284–296, 1981. URL: <https://doi.org/10.1137/0210020>, doi:10.1137/0210020.
- [7] Ahmed Bouajjani, Javier Esparza, Alain Finkel, Oded Maler, Peter Rossmanith, Bernard Willems, and Pierre Wolper. An efficient automata approach to some problems on context-free grammars. *Inf. Process. Lett.*, 74(5-

- 6):221–227, 2000. URL: [https://doi.org/10.1016/S0020-0190\(00\)00055-7](https://doi.org/10.1016/S0020-0190(00)00055-7), doi:10.1016/S0020-0190(00)00055-7.
- [8] Ahmed Bouajjani, Javier Esparza, and Oded Maler. Reachability analysis of pushdown automata: Application to model-checking. In *CONCUR '97: Concurrency Theory, 8th International Conference, Warsaw, Poland, July 1-4, 1997, Proceedings*, volume 1243 of *Lecture Notes in Computer Science*, pages 135–150. Springer, 1997.
- [9] Phillip G. Bradford. Efficient exact paths for Dyck and semi-Dyck labeled path reachability. *CoRR*, abs/1802.05239, 2018. arXiv:1802.05239.
- [10] Karl Bringmann. Personal communication. 2018.
- [11] Karl Bringmann, Allan Grønlund, and Kasper Green Larsen. A dichotomy for regular expression membership testing. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 307–318. IEEE Computer Society, 2017. URL: <https://doi.org/10.1109/FOCS.2017.36>, doi:10.1109/FOCS.2017.36.
- [12] Arnaud Carayol and Matthew Hague. Saturation algorithms for model-checking pushdown systems. In Zoltán Ésik and Zoltán Fülöp, editors, *Proceedings 14th International Conference on Automata and Formal Languages, AFL 2014, Szeged, Hungary, May 27-29, 2014*, volume 151 of *EPTCS*, pages 1–24, 2014. URL: <https://doi.org/10.4204/EPTCS.151.1>, doi:10.4204/EPTCS.151.1.
- [13] Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pages 261–270. ACM, 2016.
- [14] Krishnendu Chatterjee, Bhavya Choudhary, and Andreas Pavlogiannis. Optimal Dyck reachability for data-dependence and alias analysis. *PACMPL*, 2(POPL):30:1–30:30, 2018. URL: <https://doi.org/10.1145/3158118>, doi:10.1145/3158118.
- [15] Krishnendu Chatterjee and Georg Osang. Pushdown reachability with constant treewidth. *Inf. Process. Lett.*, 122:25–29, 2017.
- [16] Swarat Chaudhuri. Subcubic algorithms for recursive state machines. In *POPL '08*, pages 159–169. ACM, 2008.
- [17] Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *J. Symb. Comput.*, 9(3):251–280, 1990.
- [18] Mateus de Oliveira Oliveira and Michael Wehar. Intersection non-emptiness and hardness within polynomial time. In *DLT 2018*, volume 11088 of *Lecture Notes in Computer Science*, pages 282–290. Springer, 2018.
- [19] Mateus de Oliveira Oliveira and Michael Wehar. On the fine grained complexity of finite automata non-emptiness of intersection. In Natasa Jonoska

- and Dmytro Savchuk, editors, *Developments in Language Theory - 24th International Conference, DLT 2020, Tampa, FL, USA, May 11-15, 2020, Proceedings*, volume 12086 of *Lecture Notes in Computer Science*, pages 69–82. Springer, 2020. URL: https://doi.org/10.1007/978-3-030-48516-0_6, doi:10.1007/978-3-030-48516-0_6.
- [20] Danny Dolev, Shimon Even, and Richard M. Karp. On the security of ping-pong protocols. *Inf. Control.*, 55(1-3):57–68, 1982. URL: [https://doi.org/10.1016/S0019-9958\(82\)90401-6](https://doi.org/10.1016/S0019-9958(82)90401-6), doi:10.1016/S0019-9958(82)90401-6.
- [21] Henning Fernau. Modern aspects of complexity within formal languages. In Carlos Martín-Vide, Alexander Okhotin, and Dana Shapira, editors, *Language and Automata Theory and Applications - 13th International Conference, LATA 2019, St. Petersburg, Russia, March 26-29, 2019, Proceedings*, volume 11417 of *Lecture Notes in Computer Science*, pages 3–30. Springer, 2019. URL: https://doi.org/10.1007/978-3-030-13435-8_1, doi:10.1007/978-3-030-13435-8_1.
- [22] Henning Fernau and Andreas Krebs. Problems on finite automata and the exponential time hypothesis. *Algorithms*, 10(1):24, 2017. URL: <https://doi.org/10.3390/a10010024>, doi:10.3390/a10010024.
- [23] Alain Finkel, Bernard Willems, and Pierre Wolper. A direct symbolic approach to model checking pushdown systems. In Faron Moller, editor, *Second International Workshop on Verification of Infinite State Systems, Infinity 1997, Bologna, Italy, July 11-12, 1997*, volume 9 of *Electronic Notes in Theoretical Computer Science*, pages 27–37. Elsevier, 1997. URL: [https://doi.org/10.1016/S1571-0661\(05\)80426-8](https://doi.org/10.1016/S1571-0661(05)80426-8), doi:10.1016/S1571-0661(05)80426-8.
- [24] Rusins Freivalds. Fast probabilistic algorithms. In Jirí Bečvář, editor, *Mathematical Foundations of Computer Science 1979, Proceedings, 8th Symposium, Olomouc, Czechoslovakia, September 3-7, 1979*, volume 74 of *Lecture Notes in Computer Science*, pages 57–69. Springer, 1979. URL: https://doi.org/10.1007/3-540-09526-8_5, doi:10.1007/3-540-09526-8_5.
- [25] Zvi Galil. Some open problems in the theory of computation as questions about two-way deterministic pushdown automaton languages. *Mathematical Systems Theory*, 10:211–228, 1977.
- [26] Jonathan Goldstine, John K. Price, and Detlef Wotschke. A pushdown automaton or a context-free grammar: which is more economical? *Theoret. Comput. Sci.*, 18:33–40, 1982.
- [27] Jim Gray, Michael A. Harrison, and Oscar H. Ibarra. Two-way pushdown automata. *Information and Control*, 11(1/2):30–70, 1967.
- [28] Sheila A. Greibach. The hardest context-free language. *SIAM J. Comput.*, 2(4):304–310, 1973. URL: <https://doi.org/10.1137/0202025>, doi:10.1137/0202025.

- [29] Nevin Heintze and David McAllester. On the cubic bottleneck in subtyping and flow analysis. In *LICS'97*. IEEE, 1997.
- [30] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages, and Computation (3rd Edition)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2006.
- [31] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. URL: <https://doi.org/10.1006/jcss.2000.1727>, doi:10.1006/jcss.2000.1727.
- [32] Somesh Jha and Thomas W. Reps. Model checking SPKI/SDSI. *J. Comput. Secur.*, 12(3-4):317–353, 2004. URL: <http://content.iospress.com/articles/journal-of-computer-security/jcs209>.
- [33] Ivan Korec and Jirí Wiedermann. Deterministic verification of integer matrix multiplication in quadratic time. In Viliam Geffert, Bart Preneel, Branislav Rován, Julius Stuller, and A Min Tjoa, editors, *SOFSEM 2014: Theory and Practice of Computer Science - 40th International Conference on Current Trends in Theory and Practice of Computer Science, Nový Smokovec, Slovakia, January 26-29, 2014, Proceedings*, volume 8327 of *Lecture Notes in Computer Science*, pages 375–382. Springer, 2014. URL: https://doi.org/10.1007/978-3-319-04298-5_33, doi:10.1007/978-3-319-04298-5_33.
- [34] Marvin Künnemann. On nondeterministic derandomization of freivalds’ algorithm: Consequences, avenues and algorithmic progress. In Yossi Azar, Hannah Bast, and Grzegorz Herman, editors, *26th Annual European Symposium on Algorithms, ESA 2018, August 20-22, 2018, Helsinki, Finland*, volume 112 of *LIPICs*, pages 56:1–56:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. URL: <https://doi.org/10.4230/LIPICs.ESA.2018.56>, doi:10.4230/LIPICs.ESA.2018.56.
- [35] Lillian Lee. Fast context-free grammar parsing requires fast boolean matrix multiplication. *J. ACM*, 49(1):1–15, 2002.
- [36] Markus Lohrey. Algorithmics on slp-compressed strings: A survey. *Groups Complex. Cryptol.*, 4(2):241–299, 2012. URL: <https://doi.org/10.1515/gcc-2012-0016>, doi:10.1515/gcc-2012-0016.
- [37] Anders Alnor Mathiasen and Andreas Pavlogiannis. The fine-grained and parallel complexity of andersen’s pointer analysis. *Proc. ACM Program. Lang.*, 5(POPL):1–29, 2021. URL: <https://doi.org/10.1145/3434315>, doi:10.1145/3434315.
- [38] David Melski and Thomas Reps. Interconvertibility of a class of set constraints and context-free-language reachability. *Theor. Comput. Sci.*, 248(1-2):29–98, 2000.
- [39] Radford Neal. The computational complexity of taxonomic inference. Unpublished manuscript. Available at <http://www.cs.toronto.edu/~radford/ftp/taxc.pdf>, 1989.
- [40] G.C. Necula. Proof carrying code. In *POPL 97: Principles of Programming Languages*, pages 106–119. ACM, 1997.

- [41] Laurent Pierre. Rational indexes of generators of the cone of context-free languages. *Theor. Comput. Sci.*, 95(2):279–305, 1992. URL: [https://doi.org/10.1016/0304-3975\(92\)90269-L](https://doi.org/10.1016/0304-3975(92)90269-L), doi:10.1016/0304-3975(92)90269-L.
- [42] Aaron Potechin and Jeffrey O. Shallit. Lengths of words accepted by nondeterministic finite automata. *Inf. Process. Lett.*, 162:105993, 2020. URL: <https://doi.org/10.1016/j.ipl.2020.105993>, doi:10.1016/j.ipl.2020.105993.
- [43] T. Reps, S. Horwitz, and M. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *POPL 95: Principles of Programming Languages*, pages 49–61. ACM, 1995.
- [44] Wojciech Rytter. A hardest language recognized by two-way non-deterministic pushdown automata. *Inf. Process. Lett.*, 13(4/5):145–146, 1981. URL: [https://doi.org/10.1016/0020-0190\(81\)90045-4](https://doi.org/10.1016/0020-0190(81)90045-4), doi:10.1016/0020-0190(81)90045-4.
- [45] Wojciech Rytter. Fast recognition of pushdown automaton and context-free languages. *Information and Control*, 67(1-3):12–22, 1985.
- [46] Wojciech Rytter. 100 exercises in the theory of automata and formal languages, April 1987. Research report RR-99, University of Warwick, Department of Computer Science, available at <http://wrap.warwick.ac.uk/60795/>. URL: <http://wrap.warwick.ac.uk/60795/>.
- [47] Stefan Schwoon. *Model checking pushdown systems*. PhD thesis, Technical University Munich, Germany, 2002. URL: <http://tumblr.biblio.tu-muenchen.de/publ/diss/in/2002/schwoon.html>.
- [48] Joseph Swernofsky and Michael Wehar. On the complexity of intersecting regular, context-free, and tree languages. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*, volume 9135 of *Lecture Notes in Computer Science*, pages 414–426. Springer, 2015. URL: https://doi.org/10.1007/978-3-662-47666-6_33, doi:10.1007/978-3-662-47666-6_33.
- [49] Roei Tell. Proving that $prBPP=prP$ is as hard as proving that “almost NP ” is not contained in $P/poly$. *Inf. Process. Lett.*, 152, 2019. URL: <https://doi.org/10.1016/j.ipl.2019.105841>, doi:10.1016/j.ipl.2019.105841.
- [50] Leslie G. Valiant. General context-free recognition in less than cubic time. *J. Comput. Syst. Sci.*, 10(2):308–315, 1975.
- [51] Mikhail N. Vyalyi. On regular realizability problems. *Probl. Inf. Transm.*, 47(4):342–352, 2011. URL: <https://doi.org/10.1134/S003294601104003X>, doi:10.1134/S003294601104003X.

- [52] Mikhail N. Vyalyi and Alexander A. Rubtsov. On regular realizability problems for context-free languages. *Probl. Inf. Transm.*, 51(4):349–360, 2015. URL: <https://doi.org/10.1134/S0032946015040043>, doi:10.1134/S0032946015040043.
- [53] Michael Wehar. Hardness results for intersection non-emptiness. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part II*, volume 8573 of *Lecture Notes in Computer Science*, pages 354–362. Springer, 2014. URL: https://doi.org/10.1007/978-3-662-43951-7_30, doi:10.1007/978-3-662-43951-7_30.
- [54] Virginia Vassilevska Williams. Multiplying matrices faster than Coppersmith-Winograd. In *STOC*, pages 887–898. ACM, 2012.
- [55] Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity. In *International Congress of Mathematicians (ICM'18)*, 2018. Available at <https://eta.impa.br/dl/194.pdf> and <https://people.csail.mit.edu/virgi/eccentri.pdf>.
- [56] Mihalis Yannakakis. Graph-theoretic methods in database theory. In *Proceedings of the Ninth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, April 2-4, 1990, Nashville, Tennessee, USA*, pages 230–242. ACM Press, 1990.

A Proof of Claim 1

Let $\mathcal{P} = (Q, \Sigma, \Gamma, \delta, q_0, F)$. With at most a constant factor blow-up, we can assume that \mathcal{P} is in a normal form, in which each transition $(q, a, Z, q', \gamma, d) \in \delta$ is either a “push” ($\gamma = Z'Z \in \Gamma^2$) or a “pop” ($\gamma = \varepsilon$) or an “unchanged” ($\gamma = Z$).

Let $\ell = \lceil \log_2 |\Gamma| \rceil$. Fix any injective maps $\phi: \Gamma \rightarrow \{ (, [\}^\ell$ and $\psi: \Gamma \rightarrow \{),] \}^\ell$ such that, for all $Z \in \Gamma$, the word $\psi(Z)$ is obtained from $\phi(Z)$ by switching opening brackets to closing brackets without changing their type—i.e., $($ is replaced by $)$ and $[$ by $]$; and then reversing the word.

The graph G' has vertices $V \times Q \cup V' \cup \{q_f\}$, where V' is a set of new vertices and q_f is a new vertex. We shall specify V' later.

The vertices $s' = (s, q_0)$ and $t' = q_f$. There is a path from (v, q) to (v', q') labeled with the consecutive letters of $\phi(Z')$ if there is an edge $v \xrightarrow{a} v'$ in G and $(q, a, Z, q', Z'Z) \in \delta$. The intermediate vertices along this path are distinct and are not incident to any other edge of G' . Similarly, there is a path from (v, q) to (v', q') labeled with the consecutive letters of $\psi(Z)$ if there is an edge $v \xrightarrow{a} v'$ in G and $(q, a, Z, q', \varepsilon) \in \delta$. There is an edge $(v, q) \rightarrow (v', q')$ labeled with ε if there is an edge $v \xrightarrow{a} v'$ in G and $(q, a, Z, q', Z) \in \delta$. The set of all intermediate vertices added along the way constitute V' . Finally, there is an edge $(t, q) \rightarrow q_f$ labeled with ε for each $q \in F$.

Since \mathcal{P} is fixed, the algorithm runs in linear time in G and outputs G' which is linear in the size of G . By induction, we can show that there is a path from s to t in G labeled with a word from \mathcal{P} iff there is a path from (s, q_0) to q_f in G' labeled with a path in Dyck-2.

B Proof of Lemma 6

Let N be the nonterminal from the statement of the lemma. If N is not productive, i.e., cannot derive any word, then *all* of its productions can be removed without any effect on $L(\mathcal{G})$. This is simply because N cannot appear in any successful derivation. We will therefore assume that N is productive.

Consider the parse tree of any successful derivation from N . We can find in this parse tree a vertex labelled with N such that none of its descendants is labelled with N . The subtree T_N rooted at this vertex corresponds to a derivation that applies some production $P: N \rightarrow \xi$ first and never uses N again.

By removing all other productions with left-hand side N from \mathcal{G} , we obtain a new grammar \mathcal{G}' . Let us show that $L(\mathcal{G}') \neq \emptyset$. Indeed, let S be the axiom of \mathcal{G} . As S is productive, $u \in L(\mathcal{G})$ for some word u . Consider any parse tree T of u in \mathcal{G} . If T contains no occurrence of N , then it is already a valid parse tree with respect to \mathcal{G}' , and we are done. Otherwise, for every node labelled with N in T from which the shortest path to the root has no other occurrence of N , we replace the corresponding subtree by T_N . This results in a valid parse tree with respect to \mathcal{G}' , because T_N has one occurrence of N only, namely at its root, where the production applied is P . The new parse tree is a derivation of some word in $L(\mathcal{G}')$, which concludes the proof.

C Proof of Proposition 19

Let \mathcal{A} be a \mathcal{P} -automaton (saturated or not). For each $A \in \Gamma$, let $M^A = (m_{ij})$ denote the A -transition matrix of \mathcal{A} , that is, the 0–1 matrix of size $|S| \times |S|$ in which $m_{ij} = 1$ if \mathcal{A} contains a transition $i \xrightarrow{A} j$ and $m_{ij} = 0$ otherwise. Then rule (3) can be decomposed into the following updates, for all $A, B, C \in \Gamma$:

$$\begin{aligned} M^A &:= \text{bool}(M^A + T^{A,\varepsilon}), \\ M^A &:= \text{bool}(M^A + T^{A,B} \cdot M^B), \\ M^A &:= \text{bool}(M^A + T^{A,BC} \cdot M^B \cdot M^C). \end{aligned}$$

The composition of certificates (4) and the existence of verification algorithms follow as in Subsection 3.2.

D Fine-grained reductions and proof of Theorem 17

We discuss further preliminaries on fine-grained complexity, referring the reader to the recent survey by Vassilevska Williams [55] and to the paper on nondeterministic strong exponential-time hypothesis by Carmosino et al. [13].

Let L_1 and L_2 be languages, and let T_1 and T_2 be time bounds, i.e., functions $\mathbb{N} \rightarrow \mathbb{N}$. We interpret pairs (L_i, T_i) as problems with their conjectured (or presumed) complexities. We say that (L_1, T_1) *fine-grained reduces* to (L_2, T_2) , written $(L_1, T_1) \leq_{\text{FGR}} (L_2, T_2)$, if (a) for all $\varepsilon > 0$, there is $\delta > 0$ and a deterministic Turing reduction M^{L_2} from L_1 to L_2 such that $\text{DTIME}[M] \leq T_1^{1-\delta}$ and such that (b) if $Q(M, x)$ denotes the set of queries made by M to the L_2 oracle on an input x of length n , then the query lengths obey the time bound

$$\sum_{q \in Q(M, x)} (T_2(|q|))^{1-\varepsilon} \leq (T_1(n))^{1-\delta}.$$

Intuitively, a fine-grained reduction from (L_1, T_1) to (L_2, T_2) enables algorithmic savings for L_2 to be transferred to L_1 . That is, if L_2 can be solved in time $T_2^{1-\varepsilon}$, then L_1 can be solved in time $T_1^{1-\delta}$. A language L with time complexity T is **SETH-hard** if $(\text{SAT}, 2^n) \leq_{\text{FGR}} (L, T)$.

Theorem 23 ([13], Theorem 2 and Corollary 2). *Suppose NSETH holds and a problem L belongs to $\text{NTIME}[T] \cap \text{coNTIME}[T]$. Then $(\text{SAT}, 2^n) \not\leq_{\text{FGR}} (L, T^{1+\gamma})$ for any $\gamma > 0$. Also, for any L' that is SETH-hard with time T' , and any $\gamma > 0$, we have $(L', T') \not\leq_{\text{FGR}} (L, T^{1+\gamma})$.*

We are now ready to formulate Theorem 17 rigorously.

Theorem 24 (Theorem 17 restated). *Unless NSETH fails, $(\text{SAT}, 2^n) \not\leq_{\text{FGR}} (\text{D}_2\text{Reach}, n^{\omega+\gamma})$ for any $\gamma > 0$.*

It remains to observe that Theorem 17 follows from Theorem 14 and 23.

E Proof of Proposition 20

Preliminary Definitions *Two-way nondeterministic pushdown automata (2NPDA)* [27] are a powerful formalism introduced in 1967 by Gray, Harrison, and Ibarra [27].

2NPDA have the form $\mathcal{A} = (Q, \Sigma, \Gamma, \delta, q_0, F)$, where Q is a finite set of states, Σ and Γ are finite alphabets of input and stack symbols, respectively, $q_0 \in Q$ is the initial state, $F \subseteq Q$ is the set of final states, and a transition relation $\delta \subseteq Q \times \Sigma \times \Gamma \times Q \times \Gamma^* \times \{-1, 0, +1\}$. We assume Σ contains two designated “end of tape” symbols \triangleleft and \triangleright . We assume that Γ contains a designated “end of stack” symbol Z_0 such that any transition $(q, \sigma, Z_0, q', w, d) \in \delta$ satisfies $w = Z_0$. Thus, no transition of \mathcal{A} replaces Z_0 on the stack with a different symbol and no transition pushes Z_0 .

Informally, the 2NPDA \mathcal{A} has a finite control (states from Q) which reads a symbol of Σ on its input tape and the top symbol in Γ of a pushdown store. Based on the transition relation δ , 2NPDA moves by changing the control state, replacing the top symbol of the pushdown store by a finite string of symbols (possibly the empty string), and moving its input head at most one symbol left or right. Initially, the 2NPDA is in state q_0 , and its pushdown store consists of the single symbol Z_0 . The input tape consists of a word $w \in (\Sigma \setminus \{\triangleleft, \triangleright\})^*$ surrounded by a left marker \triangleleft and a right marker \triangleright and the 2NDPA scans the left marker \triangleleft .

Remark 25. We include the endmarkers \triangleleft and \triangleright in the set Σ here, even though we did not mention them back in Section 6 when specifying the alphabet Σ_0 for our hardest 2NPDA language. Naturally, all symbols used by automata (including the endmarkers) should be included in the tape alphabet of these automata.

A configuration of the 2NPDA \mathcal{A} is a triple $(q, w\hat{a}x, \gamma)$, where $q \in Q$, $w, x \in \Sigma^*$, $a \in \Sigma$, and $\gamma \in \Gamma^*$. The “hat” on a denotes that the machine is currently scanning the letter a . We write $(q_1, a_1 \dots \hat{a}_i \dots a_n, Z\gamma) \rightarrow (q_2, a_1 \dots \hat{a}_j \dots a_n, \gamma'\gamma)$ whenever $(q_1, a_i, Z, q_2, \gamma', d) \in \delta$ for $d \in \{-1, 0, +1\}$, and $j = i + d$. We require $j \in \{1, \dots, n\}$, that is, the scan position does not “fall off” the input word. Note that the input tape is not changed, only the scan position may change. We write \rightarrow^* for the reflexive and transitive closure of \rightarrow . A word $w \in (\Sigma \setminus \{\triangleleft, \triangleright\})^*$ is *accepted* by the 2NPDA if $(q_0, \triangleleft w \triangleright, Z_0) \rightarrow^* (q, \triangleleft w \triangleright, Z_0)$ for some $q \in F$. The language $L(\mathcal{A})$ of \mathcal{A} is the set of all accepted words (in $(\Sigma \setminus \{\triangleleft, \triangleright\})^*$).

Informally, the 2NPDA has some run that leads it from the initial configuration with the word on the input tape to a final state. Wlog, we can assume above that a word is accepted in a final state with the 2NPDA scanning the right end marker and the pushdown store only contains Z_0 . The transition relation is nondeterministic; we only require that some run is accepting. For the reader familiar with one-way automata, we remark that the role of epsilon-transitions is played by explicit specification of head movements.

A 1NPDA, or just PDA for short, is a 2NPDA such that $\delta \subseteq Q \times \Sigma \times \Gamma \times Q \times \Gamma^* \times \{0, +1\}$. Informally, the transitions of a PDA do not allow the scan position to move left, so PDA can only move left to right. PDA accept exactly the context-free languages. In comparison, 2NPDA are surprisingly powerful devices. In fact, even their deterministic counterparts can recognize languages such as $\{a^n b^{p(n)} \mid n \geq 0\}$ where p is a fixed polynomial with natural coefficients and $\{x\#y \mid x \text{ is a subword (factor) of } y\}$ [46, 25].

We consider the following decision problems for these machine classes. The *recognition* problem for a class of machines \mathcal{C} asks, for a fixed machine $M \in \mathcal{C}$ and an input word $w \in \Sigma^*$, if w is accepted by M , i.e., if $w \in L(M)$. The *emptiness* problem for class \mathcal{C} asks, given a machine $M \in \mathcal{C}$, if $L(M) = \emptyset$.

Proof. Proposition 20 follows from Lemmas 26, 27, and 28, which we prove next.

Lemma 26. *There exists a linear-time algorithm that, given a 2NPDA \mathcal{B} and a word w , outputs a PDA \mathcal{P} such that:*

- $|\mathcal{P}| \leq O(|w|)$ for any fixed \mathcal{B} and
- the language of \mathcal{P} is nonempty iff \mathcal{B} accepts w .

Remark. In fact, $|\mathcal{P}| \leq O(|\mathcal{B}| \cdot |w|)$.

Proof. Denote $n = |w|$ and let S be the set of control states of \mathcal{B} . Construct a PDA \mathcal{P} with the set of control states $Q = \{0, 1, \dots, n+1\} \times S$. The first component of the states of \mathcal{P} corresponds to a possible position of the input head of the 2NPDA \mathcal{B} run on w . Indeed, when \mathcal{B} is run on the word w , its head has $n+2$ possible positions: over any of the n letters of w , over the left endmarker, and over the right endmarker.

PDA \mathcal{P} has the initial state $(0, s_0)$, where s_0 is the initial state of \mathcal{B} . Transitions of the (nondeterministic) PDA \mathcal{P} are defined so that \mathcal{P} would simulate the (nondeterministic) computation of \mathcal{B} on w . The stack of \mathcal{P} is always the same as the stack of \mathcal{B} , and the second component of the control state of \mathcal{P} the same as the control state of \mathcal{B} . Transitions of \mathcal{B} depend on the input letter, which is available to \mathcal{P} , because \mathcal{P} ‘remembers’ in the control state where the input head of \mathcal{B} is positioned—and the input word w is fixed. Transitions of \mathcal{P} need not read any letter from the input; \mathcal{P} accepts (rejects) whenever so does \mathcal{B} . It is straightforward to see that both assertions of the lemma hold. \square

Lemma 27. *There exists a linear-time algorithm that, given a PDA \mathcal{P} , outputs a directed graph $G = (V, E)$, labels $\lambda: E \rightarrow \{(\cdot), [\cdot]\}$ and two vertices $s, t \in V$ such that (G, λ, s, t) is a yes-instance of $D_2\text{Reach}$ iff the language of \mathcal{P} is nonempty.*

Proof. We show how to construct the required instance of $D_2\text{Reach}$ given a PDA $\mathcal{P} = (Q, \Sigma, \Gamma, \delta, q_0, F)$.

The idea is that we encode stack symbols from Γ by sequences of words over the alphabet $\{(\cdot), [\cdot]\}$; pushing symbols on the stack corresponds to traversing edges of G labeled by opening brackets, and popping symbols—to traversing edges labeled by closing brackets.

Let $\ell = \lceil \log |\Gamma| \rceil$. Fix any injective maps $\phi: \Gamma \rightarrow \{(\cdot, [\cdot]^\ell$ and $\psi: \Gamma \rightarrow \{)\cdot,]^\ell$ such that, for all $Z \in \Gamma$, the words $\psi(Z)$ is obtained from $\phi(Z)$ by switching opening brackets to closing brackets without changing their type—i.e., $($ is replaced by $)$ and $[$ by $]$; and then reversing the word.

We next construct an auxiliary graph $G' = (V', E')$ with labels $\lambda: E \rightarrow \{(\cdot), [\cdot]\}$. The set V' contains Q as a subset. For each transition $(q, a, Z, q', \gamma, d) \in \delta$, the graph G' contains a path from q to q' of length $\ell \cdot (1 + |\gamma|)$. The edges of this path are labelled by consecutive letters of the word $\psi(Z) \cdot \phi(\gamma)$; all intermediate vertices are distinct and are incident to no other edge of G' . It is easy to see that the number of edges of G' does not exceed $|\mathcal{P}|$. (Notice that the input letter a is ignored in this construction.)

Recall that the automaton \mathcal{P} has a nonempty language if and only if there is a path from its initial configuration to a final configuration, enabled by some input word from Σ^* . The initial configuration c_0 of \mathcal{P} has control state q_0 and stack content Z_0 ; and any final configuration c has some control state $q \in F$ and the same stack content Z_0 . By construction, $c_0 \rightarrow^* c$ in the PDA \mathcal{P} if and only if the graph G' has a walk from q_0 to q labeled by a word $u \in \{(\cdot), [\cdot]\}^*$ such that $\phi(Z_0) \cdot u \cdot \psi(Z_0)$ is a Dyck-2 word.

It now remains to obtain the graph G from G' by adding fresh states s and t and connecting them to the other vertices by (1) a path from s to q_0 labeled by $\phi(Z_0)$ and

(2) paths from each $q \in F$ to t labeled by $\psi(Z_0)$. Each of these paths has length ℓ ; paths of type (2) have $\ell - 1$ edges in common. Now (G, λ, s, t) is the instance of $D_2\text{Reach}$ with the required property. \square

Lemma 28. *There exist a 2NPDA language \mathcal{L}' and a linear-time algorithm that, given a directed graph $G = (V, E)$ with labels $\lambda: E \rightarrow \{(\cdot), [\cdot]\}$ and two vertices $s, t \in V$, outputs a word w such that $w \in \mathcal{L}'$ iff (G, λ, s, t) is a yes-instance of $D_2\text{Reach}$.*

Remark 29. For the linear time bound, we assume that the graph G is encoded in binary in the input. If this is not the case, the running time of the algorithm suffers a slowdown by a factor of $O(\log |V|)$.

Proof. Words of the language \mathcal{L}' are encodings of the quadruples (G, λ, s, t) , where the vertices of G are encoded in binary. In more detail, every $w \in \mathcal{L}'$ has the following form: first an encoding of s , then an encoding of t , and finally a sequence of encodings of edges of G , where every edge $e \in E$ is followed by its label $\lambda(e)$. All these encodings are separated by delimiters.

The language \mathcal{L}' is over an alphabet of size $O(1)$; a word belongs to \mathcal{L}' iff it follows the format we have just described and the graph G has a walk from s to t labeled with a sequence from the Dyck-2 language over $\{(\cdot), [\cdot]\}$.

The algorithm from the assertion of the lemma simply writes down the encodings in the required format; it is clear that the algorithm runs in linear time and the obtained word belongs to \mathcal{L}' iff (G, λ, s, t) is a yes-instance of $D_2\text{Reach}$.

It remains to prove that the language \mathcal{L}' is recognized by a 2NPDA. Let us describe this 2NPDA \mathcal{R} . It first reads the input word and checks that it follows the format described above. If this is not the case, \mathcal{R} rejects, otherwise it guesses the required walk in G from s to t as follows.

A configuration of \mathcal{R} stores on the stack the following data:

- (at the bottom) a sequence $\sigma \in \{(\cdot), [\cdot]\}^*$, and
- (at the top) a vertex $v \in V$.

In this configuration, \mathcal{R} has already found a walk from $s \in V$ to $v \in V$ labeled with some word $\sigma' \in \{(\cdot), [\cdot]\}^*$ that reduces to σ . (A word $\sigma' \in \{(\cdot), [\cdot]\}^*$ reduces to σ if σ can be obtained from σ' by a sequence of transformations that replace the subwords $()$ and $[\]$ with ε .)

Here is how \mathcal{R} works:

1. At the beginning, initialize σ with the empty word and v with $s \in V$, pushing them to the stack.
2. Repeatedly guess the next edge $e \in E$ in the walk (leaving the loop nondeterministically after some iteration):
 - (a) move the head to the encoding of $e = (u_1, u_2)$ written on the input tape;
 - (b) pop the encoding of $v \in V$ from the stack, reading the encoding of u_1 from the input tape in sync; if $u_1 \neq v$, reject;
 - (c) look at the label $\lambda(e)$:
 - if $\lambda(e) \in \{(\cdot), [\cdot]\}$, then push $\lambda(e)$ onto the stack, extending the current $\sigma \in \{(\cdot), [\cdot]\}^*$, and

- if $\lambda(e) \in \{),]\}$, then pop the last symbol of $\sigma \in \{(, [\}^*$; proceed if the two symbols form a matching pair, otherwise reject (also reject if σ is empty);
 - (d) push the encoding of u_2 to the stack.
3. Check if the current vertex v is equal to t and σ is empty. Accept if the check succeeds, otherwise reject.

It is easy to see that an accepting computation of \mathcal{R} exists iff (G, λ, s, t) is a yes-instance of $D_2\text{Reach}$. \square

F On an $O(n^3/\log n)$ algorithm for PDA emptiness

As mentioned in Section 6, the PDA emptiness to Dyck-2 reachability reduction from Proposition 20, combined with Chaudhuri’s algorithm for CFL reachability [16], implies a slightly subcubic bound for PDA emptiness.

Indeed, Chaudhuri shows how to solve instances of CFL reachability for a fixed language (including the Dyck-2 language) in time $O(n^3/\log n)$, where n is the number of nodes in the graph.

Suppose we start with a PDA emptiness instance with s states, t transitions, and r stack symbols. Note that we can safely ignore the input alphabet symbols. The bit size of the instance is $b = O(t \log(s + r))$. The reduction from Lemma 27 gives an instance of Dyck-2 reachability with $O(b)$ nodes. Chaudhuri solves it in time $O(b^3/\log b)$, which is subcubic in the bit size of the input of PDA emptiness (although not necessarily subcubic in $s + t$).

This complexity seems folklore but was never made explicit. In particular, “textbook” algorithms for PDA emptiness go through equivalent context-free grammars [30], for which a cubic blow-up is unavoidable in the worst case [26].

G Proof of Theorem 21

Fix an arbitrary 2NPDA \mathcal{A} over a finite alphabet Σ . We can assume with no loss of generality that \mathcal{A} has a single final state and that it is different from its initial state: $|F| = 1$, $q_0 \notin F$. (It is an easy exercise to modify \mathcal{A} to ensure this assumption holds.)

Suppose an input word $w \in \Sigma^*$ is given. Lemma 26 reduces $L(\mathcal{A})$ to the emptiness problem for a PDA defined as a product of the word w and 2NPDA \mathcal{A} . More concretely, this PDA has control states $Q = \{0, 1, \dots, n+1\} \times S$ where S is the set of control states of \mathcal{A} . Note that $|Q| = O(|w| \cdot |\mathcal{A}|) = O(|w|)$ since \mathcal{A} is fixed. Here and below, the constant behind $O(\cdot)$ depends on \mathcal{A} but not on w . Similarly, the stack alphabet of the PDA is fixed too. We now give this PDA as input to a further reduction to Dyck-2 Reachability (Lemma 27), which produces an instance (G, λ, s, t) .

Claim 30. *The graph G has the following properties:*

- (a) *it has $O(|w|)$ vertices (including intermediate ones, resulting from mapping the stack alphabet into binary words);*
- (b) *its edges are labeled with symbols from $\{(,), [,]\}$;*

(c) *there is a linear order on the vertices such that each edge connects two vertices that are $O(1)$ positions away from each other in this order;*

(d) *the source is first and the sink is last in the order.*

Proof. Property (a) is due to the fact that \mathcal{A} , and thus its stack alphabet, is fixed. Property (b) is immediate. Property (c) ultimately reflects the fact that \mathcal{A} , as a two-way pushdown automaton, cannot jump cells of the input tape, that is, its head can only move one cell left or right if it moves at all —this is represented by $d \in \{-1, 0, +1\}$ in the syntax of 2NPDA. Thus, the linear order on vertices of the graph is inherited from the natural ordering of letters of the input tape, $\langle w \rangle$. Reductions to PDA emptiness and $D_2\text{Reach}$ effectively apply a direct product construction with a constant factor expansion. Within each block corresponding to an input letter, vertices can be ordered arbitrarily, provided that the initial state of \mathcal{A} comes first and the final state last —ensuring property (d). Note that our previous preprocessing of \mathcal{A} ensures that these two states are different, and our acceptance condition and subsequent reductions do the rest of the work. \square

We refer to instances (G, λ, s, t) with the properties stated in Claim 30 as those of *Restricted Dyck-2 Reachability*.

Suppose $k \in \mathbb{N}$ is chosen such that the constants behind $O(\cdot)$ in conditions (a) and (c) are at most k and every vertex has at most k outgoing edges. We think of this $k = O(1)$ as the “width” of the instance, which depends on the original 2NPDA \mathcal{A} but not on w .

Remark 31. The constant $O(1)$ in property (c) is reminiscent to the bounded pathwidth condition (see, e.g, Bienstock et al. [5]). However, in our case the graph has an even more “regular” structure. We leave it open whether this structure can be characterized by constant pathwidth *and* constant degree (and restricting the direction and labels of the edges). In comparison, Chatterjee and Osang look at pushdown reachability with constant *treewidth* [15].

It remains to map this instance of Dyck-2 Reachability to an instance of 2NPDA recognition, for a fixed 2NPDA which we now define.

For each vertex v , let $\text{index}(v)$ denote the position of v in the order specified in property (c), ranging from 1 to $O(|w|)$. (Once again, the constant behind $O(\cdot)$ depends on \mathcal{A} but not on w .) The construction below follows in spirit the proof of Lemma 28 and refines the details in order to produce a homomorphism h . The key difference is that, to produce the new input word, we will not write edges as “ $(u, v), \lambda(u, v)$ ”. Instead we will:

- 1) sort the vertices u according to their $\text{index}(u)$ ascending and, for each u , group all the edges departing from u together (each u will have at most k outgoing edges);
- 2) write edges (u, v) as pairs $(\lambda(u, v), \text{offset}(u, v))$ where $\text{offset}(u, v) = \text{index}(v) - \text{index}(u)$, i.e., how many vertices to the right the destination of the edge is; this difference is written in unary notation (without incurring blowup, as this difference cannot exceed k);
- 3) write vertices as “separators” between groups of edges.

Putting everything together, the input to the new 2NPDA has the form (5) (see page 14), where $\#$ is the vertex marker symbol, $\ell_i \in \{(\cdot), [,]\}$, and each o_i is either the empty word, or $1 \dots 1$ or $-1 \dots 1$.

Claim 32. *The set of valid encodings (5) of Restricted Dyck-2 Reachability can be recognized by a fixed 2NPDA.*

The construction of the 2NPDA in Claim 32 is similar to the reduction of Lemma 28 which we already have. Instead of guessing the next vertex, this new 2NPDA \mathcal{A}'_0 “scrolls” left and right in a deterministic way to the destination of the current edge, counting in unary with the help of its stack. The nondeterministic choices that \mathcal{A}'_0 makes are which outgoing edge from the current vertex to choose next.

Note that the construction of \mathcal{A}'_0 is independent of k , thus identifying a single hardest language, $L(\mathcal{A}'_0)$. Moreover, for a given initial 2NPDA \mathcal{A} this reduction replaces each symbol in w with $O(1)$ vertices and $O(1)$ edges, where this $O(1)$ depends just on \mathcal{A} and not w . The exact collection of these vertices and edges is fully determined by each symbol of w , independently of its position within w . The vertices are not addressed in any “absolute” numbering scheme — so this mapping can be realised as a homomorphism.

Remark 33. The use of relative rather than absolute addresses (to encode offsets) appears in a related context but for a different problem in Neal’s work on taxonomic inference [39], which is at the origin of the connection between 2NPDA and program analysis.

Summary and the endmarkers problem. We now have achieved the following: for every 2NPDA \mathcal{A} there is a homomorphism h_1 such that $w \in L(\mathcal{A})$ if and only if $h_1(\triangleleft w \triangleright) \in L(\mathcal{A}'_0)$. Note the appearance of the endmarkers here. (We use \triangleleft instead of \triangleleft and \triangleright instead of \triangleright to avoid a notation clash in the discussion that follows.) They reflect the fact that, in the chain of our reductions, the set of control states of the PDA is $\{0, 1, \dots, n + 1\} \times S$ not $\{1, \dots, n\} \times S$.

To lift our construction from $\triangleleft w \triangleright$ to just w , it may be tempting to appeal to the following fact, which is not difficult to prove. Let $x, y \in \Sigma^*$ be fixed. Suppose a 2NPDA accepts a language $L \subseteq x \cdot \Sigma^* \cdot y$. Then there exists another 2NPDA which accepts the language $\{w \mid xwy \in L\}$.

Unfortunately, this fact does not quite achieve our goal. This is because the new 2NPDA we would obtain from it depends on x and y . In our context, x and y should be the images of the original endmarkers, i.e., we would like to have $x = h_1(\triangleleft)$ and $y = h_1(\triangleright)$. But these two words depend on the homomorphism h_1 , and thus on the 2NPDA \mathcal{A} that we started from. This is at odds with our objective: we need a single 2NPDA for our hardest language, not an entire family dependent on \mathcal{A} .

There are several ways to deal with this issue. One is reminiscent of Rytter’s approach [44]: we can decide we are content with keeping a single endmarker in, i.e., we would only like to find an L_0 such that, for all $w \in \Sigma^+$, one has $w \in L$ iff $h(w\$) \in L_0$. Here $\$$ is a fresh symbol. It is not very difficult to find such an h and L_0 based on our construction: essentially, the word $h_1(\triangleleft)$ needs to be merged with the word $h_1(\triangleright)$ and placed to the right of $h_1(w)$. So we would like to choose $h(\$) = h_1(\triangleright)h_1(\triangleleft)$ and $h(a) = h_1(a)$ for all other symbols a . The 2NPDA for L_0 is the same as our 2NPDA \mathcal{A}'_0 constructed above, with the following modification. Suppose it starts following an edge from some vertex (block) to the left but hits the left end of the tape, i.e., the left endmarker \triangleleft . We now use this symbol to refer to the tape alphabet of the 2NPDA \mathcal{A}'_0 (and not the tape alphabet of the original machine \mathcal{A}). The new 2NPDA will move all the way to the right end of the tape and continue its search for the destination vertex from the right endmarker \triangleright . Edges within $h_1(\triangleright)$ need not be changed, but the ones among them

that lead to the right ($\text{offset}(u, v) > 0$, or equivalently $o_i \in 1^+$) will make the 2NPDA hit \triangleright , go back to the left of the tape and continue the seart from \triangleleft .

One further technicality that needs to be dealt with is the beginning and end of the computation. Recall that our Restricted Dyck-2 Reachability asked for a path from the very first vertex to the very last one. Since $h_1(\trianglelefteq)$ moved, we now need to change this convention. More concretely, the only two vertices that we can distinguish correspond to the *last two* control states and the head position over the *left* endmarker. So the original 2NPDA \mathcal{A} needs to be changed accordingly.

While this approach recovers Rytter’s result, we show below that there is a way to eliminate the extra symbol $\$$ altogether.

Merging endmarker blocks into other symbols. Our solution to the endmarkers problem acknowledges that the words $h_1(\trianglelefteq)$ and $h_1(\triangleright)$ cannot be eliminated completely. Indeed, the vertices and edges that these two words encode correspond to the behaviour of the original 2NPDA \mathcal{A} over the tape endmarkers, and this behaviour can contribute to the computations of \mathcal{A} in a nontrivial way.

However, what we can do is to embed all this information into words $h(a)$ for *all* other symbols a . For a first intuition (to be amended later), we would like to set $h(a) = h_1(\trianglelefteq) \parallel h_1(a) \parallel h_1(\triangleright)$ for all non-endmarker symbols a , where \parallel denotes a specially tailored ternary version of the perfect shuffle operation. More concretely, let w_1, w_2, w_3 be arbitrary words such that, for some single ℓ , we have $w_i = \prod_{j=1}^{\ell} \#w_{i,j}$ where none of the words $w_{i,j}$ contains the vertex marker symbol $\#$. Then

$$w_1 \parallel w_2 \parallel w_3 := \prod_{j=1}^{\ell} \#w_{1,j} \#w_{2,j} \#w_{3,j} .$$

Note that this shuffling relies on ℓ being the same for all three arguments, and ultimately this means the same number of vertices (blocks) in all words $h_1(a)$. This is in fact ensured by our constructions above (although we could always achieve this by introducing extra dummy vertices where necessary).

As a result of this shuffling arrangement, we can think of new input words as having three interleaving “tracks”, each containing a separate sequence of vertices. Naturally, this requires some changes to the wiring, as follows.

First, the offsets that specify the edges of the graph departing from the vertices of $h_1(a)$ need to be updated. This is not difficult. Recall that edge destinations are specified using relative addresses of vertices. For every edge from a vertex in $h_1(a)$, its offset needs to be multiplied by 3, so that the edge skips intermediate vertices from copies of $h_1(\trianglelefteq)$ and $h_1(\triangleright)$.

Second, we need to provide a way for the new 2NPDA \mathcal{A}_0 to reach the vertices in $h_1(\trianglelefteq)$ and $h_1(\triangleright)$. To achieve this, we consider the scenario in which \mathcal{A}_0 will traverse edges leading to a vertex in $h_1(\trianglelefteq)$. (The case of $h_1(\triangleright)$ is handled in a symmetric way.) Suppose the head of the 2NPDA is over a block (vertex) within the leftmost $h_1(a)$. Taking an edge with a negative offset, it moves left but then hits the left tape endmarker \triangleleft . When it does so, the stack of the 2NPDA still contains the number of vertices to be skipped. The 2NPDA then needs to change from the second (main) track, which contains the information from $h_1(a)$ s, to the first track, which stores multiple copies of the word $h_1(\trianglelefteq)$. Effectively, this amounts to treating \triangleleft as just another $\#$ that on top of its usual function makes the machine change direction. After that, however, we see that the number of

vertices to be skipped was counted from the right of $h_1(\trianglelefteq)$ and not from the left where the head of the automaton is now located. Thus, we *reverse* the encoding of each of $h_1(\trianglelefteq)$ and $h_1(\trianglerighteq)$, as follows: we re-define our special shuffle as

$$w_1 \parallel w_2 \parallel w_3 := \prod_{j=1}^{\ell} \# \bar{w}_{1,\ell+1-j} \# w_{2,j} \# \bar{w}_{3,\ell+1-j} ,$$

where \bar{u} is the same word as u in which every maximal subword of the form -1^m is replaced with 1^m and each 1^m , without a preceding $-$, with -1^m . Our 2NPDA must remember, in its control state, which “track” of the input it is over. The second track corresponds to the usual operation. Over the first track:

- Edges previously specified by positive offset needs to followed to the *left* instead of to the right (hence the $\bar{w}_{1,\ell+1-j}$ above and not $w_{1,\ell+1-j}$. If the left tape endmarker \triangleleft is encountered, the automaton transitions to the second (main) track, and only then continues to the right (in the normal mode).
- Edges specified by the negative offset need to be followed to the *right* instead of to the left (again, this matches the $\bar{w}_{1,\ell+1-j}$ above). We note that if the input word for our 2NPDA is the homomorphic image under h of some word in Σ^* , then the automaton will never leave the leftmost $h(a)$ while being on the first track, because the original 2NPDA \mathcal{A} cannot move left from the left endmarker.

The third track is arranged in a similar way.

Importantly, while we apply these changes to h and the “wiring” of the graph, we can keep the semantics of our hardest language untouched. The “tracks” themselves need not enter the description of the language. The only new “feature” that is necessary is *changing* the tracks — and this can be achieved simply by specifying that when our new 2NPDA \mathcal{A}_0 encounters a tape endmarker during its operation, this endmarker is counted as a virtual vertex and “reflects” off it, continuing the countdown in the opposite direction.

However, as was the case with the approach described above and involving $\$,$ our new construction of h breaks the convention about the source and target vertices in the Dyck-2 reachability instance (albeit in a slightly different way). Because of the effective reversal of vertex ordering within $h_1(\trianglelefteq)$ and $h_1(\trianglerighteq)$, the required adjustment to the original 2NPDA \mathcal{A} is that its initial control state needs to be the *last* and its (only) final control state the *first* in the ordering.

To sum up, by applying these adjustments and “compiling” the homomorphism h the way we have described, we arrive at the desired 2NPDA \mathcal{A}_0 . (Note that there is freedom in whether we take $\varepsilon \in L(\mathcal{A}_0)$ or $\varepsilon \notin L(\mathcal{A}_0)$, but as some 2NPDA languages contain ε and some do not, their homomorphic images will necessarily disagree on ε , no matter our choice of the homomorphism.) This completes the proof.