Electronic Colloquium on Computational Complexity, Report No. 16 (2020)

Hitting Sets Give Two-Sided Derandomization of Small Space

Kuan Cheng*
Dept. of Computer Science
University of Texas at Austin
ckkcdh@hotmail.com

William M. Hoza†
Dept. of Computer Science
University of Texas at Austin
whoza@utexas.edu

Abstract

A *hitting set* is a "one-sided" variant of a pseudorandom generator (PRG), naturally suited to derandomizing algorithms that have one-sided error. We study the problem of using a given hitting set to derandomize algorithms that have *two-sided* error, focusing on space-bounded algorithms. For our first result, we show that if there is a log-space hitting set for polynomial-width read-once branching programs (ROBPs), then not only does $\mathbf{L} = \mathbf{RL}$, but $\mathbf{L} = \mathbf{BPL}$ as well. This answers a question raised by Hoza and Zuckerman [HZ18].

Next, we consider constant-width ROBPs. We show that if there are log-space hitting sets for constant-width ROBPs, then given black-box access to a constant-width ROBP f, it is possible to deterministically estimate $\mathbb{E}[f]$ to within $\pm\varepsilon$ in space $O(\log(n/\varepsilon))$. Unconditionally, we give a deterministic algorithm for this problem with space complexity $O(\log^2 n + \log(1/\varepsilon))$, slightly improving over previous work.

Finally, we investigate the limits of this line of work. Perhaps the strongest reduction along these lines one could hope for would say that for every explicit hitting set, there is an explicit PRG with similar parameters. In the setting of constant-width ROBPs over a large alphabet, we prove that establishing such a strong reduction is at least as difficult as constructing a good PRG outright. Quantitatively, we prove that if the strong reduction holds, then for every constant $\alpha > 0$, there is an explicit PRG for constant-width ROBPs with seed length $O(\log^{1+\alpha} n)$. Along the way, unconditionally, we construct an improved hitting set for ROBPs over a large alphabet.

*Supported by a Simons Investigator Award (#409864, David Zuckerman).
†Supported by the NSF GRFP under Grant DGE-1610403 and by a Harrington Fellowship from UT Austin.

ISSN 1433-8092

1 Introduction

Suppose some decision problem can be solved by an efficient randomized algorithm. That's good, but an efficient deterministic algorithm would be even better. We would therefore like to deterministically analyze the acceptance probability of the randomized algorithm on a given input. An ambitious approach to derandomization is to try to design a suitable *pseudorandom generator* (PRG).

Definition 1.1. *Let $\mathcal{F}$ be a class of functions $f\colon \{0,1\}^n \to \{0,1\}$. An ε-*PRG *for $\mathcal{F}$ is a function $G\colon \{0,1\}^s \to \{0,1\}^n$ such that for every $f \in \mathcal{F}$, $\left|\mathbb{E}[f] - \mathbb{E}_{X\in\{0,1\}^s}[f(G(X))]\right| \leq \varepsilon$.*

Let n be the number of random bits used by the randomized algorithm, and ensure that $\mathcal{F}$ can compute the action of the randomized algorithm on its random bits. By iterating over all "seeds" $x \in \{0,1\}^s$ and plugging $G(x)$ into the randomized algorithm, we can get an estimate of its acceptance probability with additive error ε.

Unfortunately, designing efficient PRGs has proved to be extremely difficult. Constructing a *hitting set* is sometimes less difficult.

Definition 1.2. *Let $\mathcal{F}$ be a class of functions $f\colon \{0,1\}^n \to \{0,1\}$. An ε-*hitting set *for $\mathcal{F}$ is a set $H \subseteq \{0,1\}^n$ such that for every $f \in \mathcal{F}$ with $\mathbb{E}[f] \geq \varepsilon$, there is some $x \in H$ such that $f(x) = 1$.*

The image of any PRG is clearly a hitting set. By iterating over all strings in a hitting set, we can at least distinguish acceptance probability 0 from acceptance probability $\geq \varepsilon$. This is already sufficient for derandomizing some algorithms (namely, those with "one-sided error"). In this paper, we investigate the possibility of using a hitting set in a nontrivial way to obtain an estimate of the acceptance probability with a small additive error, just like what a PRG would have provided.

This possibility was previously studied in the context of derandomizing time-bounded algorithms. Several proofs have been discovered showing that if there is a polynomial-time hitting set for size-n circuits, then $\mathbf{P} = \mathbf{BPP}$ [ACR96, BF99, ACRT99, GVW11]. In Appendix A we provide yet another proof of this theorem; our short proof is arguably simpler than all previous proofs. However, the focus of our paper is derandomizing space-bounded algorithms.

1.1 Derandomizing log-space algorithms

The behavior of a small-space algorithm as a function of its random bits can be modeled by a *read-once*[1] *branching program* (ROBP). A width-w length-n ROBP is a directed graph consisting of $n+1$ layers with w vertices per layer. There is a designated "start vertex" in the first layer. Every vertex not in the last layer has two outgoing edges labeled 0 and 1 leading to the next layer. An n-bit input string naturally identifies a path through the graph by reading from left to right. The program accepts or rejects this string depending on whether the path ends at the designated "accept vertex" in the last layer.

Recall that **BPL** and **RL** are the classes of languages that can be decided by randomized log-space algorithms that always halt with two-sided and one-sided error respectively. A log-space hitting set for polynomial-width ROBPs would immediately imply $\mathbf{L} = \mathbf{RL}$. For our first result, we show that such a hitting set would also imply $\mathbf{L} = \mathbf{BPL}$.

Theorem 1.3. *Assume that for every $n \in \mathbb{N}$, there is a $\frac{1}{2}$-hitting set for width-n, length-n ROBPs that can be computed in space $O(\log n)$. Then $\mathbf{L} = \mathbf{BPL}$.*

[1] Because space-bounded algorithms only have read-once access to their random bits, it does not seem possible to adapt the existing derandomizations of **BPP** using a hitting set to the **BPL** case.

1.2 Motivation: Recent work on hitting sets

Theorem 1.3 is especially interesting in light of recent constructions of improved hitting sets for ROBPs [BCG18, HZ18]. The best known PRG for polynomial-width ROBPs is still Nisan's PRG [Nis92], which has seed length

$$O(\log^2 n + \log n \log(1/\varepsilon)).$$

Until recently, Nisan's PRG also provided the best hitting set for polynomial-width ROBPs. Using sophisticated and novel techniques, Braverman, Cohen, and Garg obtained a hitting set with space complexity

$$\widetilde{O}(\log^2 n + \log(1/\varepsilon)),$$

which is an improvement when ε is very small [BCG18].

Actually, Braverman, Cohen, and Garg constructed something better than a hitting set, called a *pseudorandom pseudodistribution* (PRPD).

Definition 1.4 ([BCG18])**.** *Let $\mathcal{F}$ be a class of functions $f\colon \{0,1\}^n \to \{0,1\}$. An* ε-PRPD *for $\mathcal{F}$ is a function $D\colon \{0,1\}^n \to \mathbb{R}$ such that for every $f \in \mathcal{F}$,*

$$\left| \sum_{x \in \{0,1\}^n} f(x)D(x) - \mathbb{E}[f] \right| \leq \varepsilon.$$

A PRPD can be used to estimate $\mathbb{E}[f]$ to within $\pm\varepsilon$, provided there is an efficient algorithm that enumerates all $x \in \mathrm{supp}(D)$ and computes $D(x)$. The concept of a PRPD generalizes the concept of a PRG, because given a PRG G with seed length s, one can set $D(x) = |G^{-1}(x)| \cdot 2^{-s}$. In turn, if D is a PRPD, then $\mathrm{supp}(D)$ is a hitting set. So a PRPD is *intermediate* between a hitting set and a genuine PRG.

After Braverman, Cohen, and Garg's work [BCG18], Hoza and Zuckerman gave a simpler construction of an ε-hitting set for polynomial-width ROBPs, with the slightly improved seed length $O(\log^2 n + \log(1/\varepsilon))$ [HZ18]. Their construction is weaker in that it does not provide a PRPD. Theorem 1.3 bridges the gap between the two concepts somewhat: by Theorem 1.3, any generic hitting set can be used for two-sided derandomization, which was the main strength of a PRPD over a hitting set in the first place.

1.3 The constant-width setting

However, there is a weakness of Theorem 1.3. A PRG or a PRPD would provide a *black-box* derandomization, whereas the algorithm of Theorem 1.3 is not black-box. This weakness is especially acute when we consider the constant-width case. Given a constant-width ROBP f directly as input, it is trivial to compute $\mathbb{E}[f]$ with high accuracy, so the algorithm of Theorem 1.3 is meaningless. Nevertheless, constant-width ROBPs can compute many interesting functions, and it is a major open challenge to design improved PRGs, PRPDs, or hitting sets for constant-width ROBPs. (For width 2, optimal PRGs are known [BDVY13]. For width 3, the current best PRG has seed length $\widetilde{O}(\log n \log(1/\varepsilon))$ [MRT19]. The best hitting sets for width 3 are superior, with space complexity $\widetilde{O}(\log(n/\varepsilon))$ for small ε [GMR+12] or $O(\log n)$ for $\varepsilon \approx 1$ [ŠŽ11]. For width 4, the state of the art is simply the best results for polynomial-width ROBPs.)

To address this weakness of Theorem 1.3, we abstract the "black-box" feature of PRGs and PRPDs in the following definition.

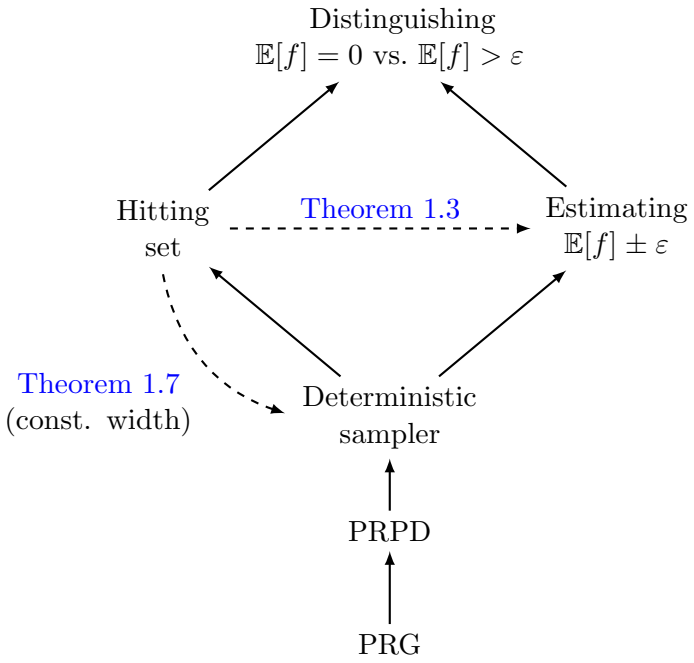

Figure 1: The relationships between different derandomization goals. The solid arrows are implications that are immediate from the definitions and hold for essentially any class $\mathcal{F}$. The dashed arrows are theorems in this paper, holding for ROBPs specifically.

Definition 1.5. *Let $\mathcal{F}$ be a class of functions $f\colon \{0,1\}^n \to \{0,1\}$. A* deterministic ε-sampler *for $\mathcal{F}$ is a deterministic oracle algorithm A that outputs a real number such that for every $f \in \mathcal{F}$,*

$$|A^f - \mathbb{E}[f]| \leq \varepsilon.$$

The concept of a deterministic sampler generalizes that of a PRPD, because given a PRPD D, one can set $A^f = \sum_x f(x)D(x)$. In the other direction, deterministic samplers imply hitting sets.

Proposition 1.6. *Identify 0 with the constant 0 function on $\{0,1\}^n$, and assume $0 \in \mathcal{F}$. Let A be a deterministic ε-sampler for $\mathcal{F}$, and let $H \subseteq \{0,1\}^n$ be the set of points where A^0 queries its oracle. Then for every $\varepsilon' > 2\varepsilon$, H is an ε'-hitting set for $\mathcal{F}$.*

Proof. Let $f \in \mathcal{F}$ satisfy $\mathbb{E}[f] > 2\varepsilon$. Since $|A^0 - 0| \leq \varepsilon$ and $|A^f - \mathbb{E}[f]| \leq \varepsilon$, $A^0 \neq A^f$. Therefore, A^f must query f at some point $x \in f^{-1}(1)$. The *first* such query must be at a point $x \in H$. □

All known derandomizations of **BPP** using a hitting set [ACR96, BF99, ACRT99, GVW11], including our new derandomization in Appendix A, are black-box. That is, one can generically "upgrade" a polynomial-time hitting set for size-n circuits into a polynomial-time deterministic sampler for size-n circuits. For our second result, we prove the analogous reduction for constant-width ROBPs. (See Figure 1.)

Theorem 1.7. *Assume that for every constant w, for all $n \in \mathbb{N}$, there is a $\frac{1}{2}$-hitting set for width-w length-n ROBPs that can be computed in space $O(\log n)$. Then for every constant w, for all $n \in \mathbb{N}$ and all $\varepsilon > 0$, there is a deterministic ε-sampler for width-w length-n ROBPs that runs in space $O(\log(n/\varepsilon))$.*

The proof of Theorem 1.7 uses different techniques than that of Theorem 1.3. The space complexity of our deterministic sampler is proportional to the width parameter w (see Theorem 4.1), so the sampler becomes meaningless when w is large. Thus, Theorems 1.3 and 1.7 are incomparable.

We also obtain a new *unconditional* deterministic sampler. Using prior work, the best deterministic sampler for constant-width ROBPs was from Braverman, Cohen, and Garg's PRPD [BCG18] if ε is small (space complexity $\widetilde{O}(\log^2 n + \log(1/\varepsilon))$, or else just from Nisan's PRG [Nis92] if ε is not so small (space complexity $O(\log^2 n + \log n \log(1/\varepsilon))$). By applying the reduction underlying Theorem 1.7 to the hitting set of Hoza and Zuckerman [HZ18], we achieve a slight improvement.

Theorem 1.8 (Unconditional sampler)**.** *For every constant w, for all $n \in \mathbb{N}$ and all $\varepsilon > 0$, there is a deterministic ε-sampler for width-w length-n ROBPs running in space $O(\log^2 n + \log(1/\varepsilon))$.*

In light of Theorem 1.8, when it comes to deterministic samplers, there is now a slight gap between the state of the art for polynomial-width ROBPs vs. the state of the art for width-w ROBPs with w a large constant. In other words, Theorem 1.8 is a case where we can take advantage of narrowness. There is no such gap when it comes to PRGs, PRPDs, or hitting sets.

1.4 Negative result

Theorem 1.7 raises the question of whether we can go even further and upgrade any hitting set into a genuine PRG. In the time-bounded setting, this is indeed possible via the "hardness vs. randomness" paradigm. (If for every n there is a hitting set for size-n circuits computable in $\mathrm{poly}(n)$ time, then there is a language in $\mathbf{E}$ that requires circuits of size $2^{\Omega(n)}$. A major achievement in complexity theory was to show that assuming such a language exists, for every n, there is a polynomial-time logarithmic-seed PRG for size-n circuits [IW97].) Also, in the context of low-degree polynomials, Bogdanov showed how to convert any hitting set with a certain density property into a PRG [Bog05]. Can a similar reduction be proven for small-space models?

We focus on the setting of constant-width ROBPs over a large alphabet. (An *ROBP over the alphabet* Σ computes a function $f\colon \Sigma^n \to \{0,1\}$; each vertex not in the last layer has $|\Sigma|$ outgoing edges labeled with the symbols in Σ.) We prove that *if* for every explicit hitting set in this setting, there is an explicit PRG with similar parameters, then there is in fact an explicit PRG for constant-width binary ROBPs with seed length $O(\log^{1+\alpha} n)$, where $\alpha > 0$ is an arbitrarily small constant. See Theorem 5.3 for the precise statement.

1.4.1 Interpretation

Like any conditional theorem, Theorem 5.3 has both a positive and a negative interpretation.[2] According to the negative interpretation, Theorem 5.3 shows that it would be difficult to establish a general reduction from PRGs to hitting sets. After all, it's as difficult as constructing a good PRG for constant-width ROBPs, which is a challenge that researchers have been struggling with for decades. In this sense, Theorem 5.3 provides an "excuse" for the fact that Theorems 1.3 and 1.7 do not provide genuine PRGs.

We feel that the negative interpretation is more realistic, but there is also a sensible positive interpretation. According to the positive interpretation, our work provides a new approach to constructing improved PRGs or hitting sets for constant-width ROBPs. One "merely" needs to bridge the gap between deterministic samplers and PRGs. This could be done in one of two ways. One could improve Theorem 1.7 so that it concludes with a PRG instead of a deterministic sampler.

[2] Throughout this discussion, we will ignore the issue of alphabet size, to simplify matters. The proof of Theorem 1.7 does generalize well to the large-alphabet case.

Alternatively, one could improve the construction of Theorem 5.3 so that rather than relying on the equivalence of hitting sets and PRGs, it merely relies on the equivalence of hitting sets and deterministic samplers. (In exchange, presumably the conclusion would merely be a deterministic sampler rather than a true PRG, but that would still be a breakthrough.)

1.5 Overview of techniques

1.5.1 Techniques for Theorem 1.3

We begin by outlining the proof of Theorem 1.3 (on derandomizing **BPL**). Suppose we are given as input a width-n length-n ROBP f. To derandomize **BPL**, it suffices to estimate $\mathbb{E}[f]$ to within a small additive error. We use the given hitting set $H \subseteq \{0,1\}^{\text{poly}(n)}$. We think of a string $x \in H$ as providing, for each vertex v in f, a list of $\text{poly}(n)$ "sample inputs," and we compute the fraction of those sample inputs that lead to v.

That fraction is an estimate of the probability that a random walk from the start vertex of f reaches v. If x were chosen at random, with high probability, each estimate would be a close approximation to the truth. However, for $x \in H$, the estimates are not necessarily good. We look for an $x \in H$ such that these estimates are at least *locally consistent*, i.e., the estimates for each pair of consecutive layers are consistent with the arrangement of edges between those layers. Having found such an $x \in H$, we output the corresponding estimated probability of reaching the accept vertex of f.

A straightforward calculation shows that if the estimates are locally consistent with one another, then each estimate is indeed close to the corresponding true probability. To complete the proof, we must use the fact that H is a hitting set to show that there is always some $x \in H$ that passes the local consistency test. The local consistency test can easily be computed in small space, but that involves reading the bits of x multiple times, so H is not immediately guaranteed to hit it.

Instead, we observe that there exists a polynomial-width ROBP that reads x and determines whether each estimate is close to the corresponding *true* probability. The ROBP simply has the true probabilities hard-coded in. There is no need to algorithmically construct that ROBP; the mere fact that it exists implies the existence of an $x \in H$ such that the estimates are all close to the corresponding true probabilities. This readily implies that the estimates are locally consistent with one another.

1.5.2 Techniques for Theorem 1.7

The proof of Theorem 1.7 (on deterministic samplers) uses different techniques. Let f be a constant-width ROBP. To estimate $\mathbb{E}[f]$, we attempt to work our way backward through the branching program, computing the probabilities of acceptance from each vertex. This plan is complicated by the fact that we only have black-box access to f. At a high level, for each layer, we use the assumed hitting set H to approximately compute the transitions at that layer, which allows us to continue computing the probabilities of acceptance from each vertex.

In more detail, the hitting set assists us in two different ways. First, we identify each prefix of a string in H with the vertex that is reached when f reads the prefix. In this way we are able to "find" all the vertices of f – or at least, all non-negligible vertices.

However, we are now effectively dealing with a width-$|H|$ branching program, because we have a copy of v for each string in H that leads to v. This interferes with our plan, because $|H| = \text{poly}(n)$ and hence we cannot afford to store the acceptance probabilities of all vertices in a single layer. The second way we use H is to determine which of these vertices are redundant. If there is some string in H that leads to accept from one vertex and reject from another, then the two vertices are

not equivalent. Otherwise, the two vertices can be safely merged, because they must be two copies of the same vertex in f – or at least, they must correspond to two very similar vertices in f. The merging condition can be checked by making queries to f. By merging vertices, we effectively bring the width back down to a constant.

Unfortunately, the fact that two vertices are equivalent does not imply that their outneighbors are equivalent, so it is not immediately clear how to "merge" the outgoing edges. We show that it suffices to retain the outgoing edges from whichever vertex has the higher acceptance probability.

1.5.3 Techniques for Theorem 5.3

Recall that to prove Theorem 5.3, we must (conditionally) construct a PRG with seed length $O(\log^{1+\alpha} n)$, where $\alpha > 0$ is an arbitrarily small constant. For simplicity, in this overview, we will focus on the case $\alpha = 1/2$, i.e., seed length $O(\log^{3/2} n)$. Recall also that we are focusing on the constant-width case.

The starting point of the construction is the INW PRG, which ε-fools constant-width ROBPs over the alphabet $\{0,1\}^t$ with seed length $O(t + \log(n/\varepsilon)\log n)$ [INW94]. (Nisan's PRG [Nis92] does not achieve the same optimal dependence on t.) Next, we present a reduction, showing how to convert a PRG with moderate error into a hitting set with very small threshold (Theorem 5.4). Hoza and Zuckerman gave a similar reduction [HZ18], but their reduction only applies to binary ROBPs (the case $t = 1$). Our reduction is based on a more sophisticated variant of a key lemma in Hoza and Zuckerman's work [HZ18].

Applying our new reduction to the INW generator, we unconditionally obtain an improved hitting set. The best previous hitting sets had space complexity $O(t + \log^2 n + \log(1/\varepsilon)\log n)$ [INW94] or $O(t \log n + \log^2 n + \log(1/\varepsilon))$ [HZ18]. Our new hitting set (Corollary 5.8) achieves the "best of both worlds," with space complexity $O(t + \log^2 n + \log(1/\varepsilon))$.

The next step in the proof of Theorem 5.3 is to apply the assumption of Theorem 5.3, converting our hitting set into a PRG. The final step is to use traditional "seed recycling" techniques to trade the excellent dependence on ε for an improved dependence on n. Briefly, starting with a length-n ROBP over the alphabet $\{0,1\}^t$, we first use a randomized sampler [GW97] to reduce the alphabet size to $\text{poly}(n)$. Then we divide our length-n ROBP of interest into blocks of length $m = 2^{\sqrt{\log n}}$. We can fool each chunk to within error $1/\text{poly}(n)$ using a seed of length $O(\log^2 m + \log n) = O(\log n)$. Using the randomized sampler again, this allows us to effectively pay $O(\log n)$ truly random bits and reduce the length of the branching program by a factor of m. After repeating this process $\sqrt{\log n}$ times, the length is reduced to a constant, and we have paid a total of $O(\log^{3/2} n)$ truly random bits. (To achieve seed length $O(\log^{1+\alpha} n)$, we start the whole process over again and iterate roughly $1/\alpha$ times.)

1.6 Related work

We have already referenced most of the work related to this paper, such as work on derandomizing **BPP** using a hitting set [ACR96, BF99, ACRT99, GVW11]. However, a few additional papers deserve mention.

1.6.1 BPL $\subseteq$ ZP*L

Our derandomization of **BPL** given a hitting set is similar to Nisan's unconditional proof that **BPL** $\subseteq$ **ZP*****L** [Nis93]. To estimate the acceptance probability of a width-n length-n ROBP f, Nisan, like us, interprets a string $x \in \{0,1\}^{\text{poly}(n)}$ as a list of sample inputs, which he uses to compute estimates for the probabilities of reaching each vertex of f. Nisan's algorithm picks x

at random, and then in a similar fashion as our algorithm, performs certain "local tests" at each vertex to verify that the sample inputs are trustworthy. Nisan's local tests can be computed in small space given two-way access to x, and passing the local tests implies that the estimates are close to the corresponding true probabilities. Our local consistency test also satisfies these properties, and indeed, one can obtain an alternative proof that $\mathbf{BPL} \subseteq \mathbf{ZP}^*\mathbf{L}$ from our analysis. However, a technical point is that we use *fresh samples* for each vertex, whereas Nisan uses one set of n-bit sample inputs for all the vertices. This crucial distinction is how we are able to ensure the existence of a polynomial-width ROBP that compares our estimates to the true probabilities. Unfortunately, using fresh samples breaks Nisan's local tests, hence our new local consistency test.

1.6.2 Deterministically simulating BPL with very low error

The current best hitting sets for polynomial-width ROBPs [BCG18, HZ18] are superior to the best known PRGs [Nis92] when ε is very small. One might hope that by plugging in the recent hitting sets, our reductions could provide a new unconditional deterministic algorithm for estimating the acceptance probability of a **BPL** algorithm to within $\pm\varepsilon$, with an improved space complexity when ε is very small. Unfortunately, this idea doesn't get off the ground, because to estimate the acceptance probability to within $\pm\varepsilon$, we rely on a $\frac{1}{2}$-hitting set for ROBPs of length $\mathrm{poly}(n/\varepsilon)$ rather than an ε-hitting set for ROBPs of length n. The good news is that Ahmadinejad et al. recently tackled this same problem with different techniques. They designed an algorithm that runs in space $O(\log^{3/2} n + \log n \log\log(1/\varepsilon))$ [AKM+19].

1.6.3 Using equivalences to conditionally construct PRGs

Our theorem about the limitations of this line of work (Theorem 5.3) is similar to a result by Hoza and Umans [HU17]. Like us, Hoza and Umans showed that if PRGs are equivalent to a seemingly weaker notion, then the equivalence itself can be used to construct a good PRG. Hoza and Umans focused on the distinction between PRGs and non-black-box derandomization, whereas we focus on the distinction between PRGs and hitting sets. The proofs of the two theorems rely on similar iterative strategies, but the specific reductions are different.

1.7 Outline of this paper

In Section 3, we present our derandomization of **BPL** given a hitting set for polynomial-width ROBPs. In Section 4, we present our deterministic sampler for constant-width ROBPs given a hitting set. Finally, in Section 5, we present our theorem on the limitations of this line of work.

2 Notation

Let U_n denote the uniform distribution over $\{0,1\}^n$. For two strings x, y, let $x \circ y$ denote the concatenation of x with y. Suppose an ROBP f is clear from context. Let v_{start} denote the start vertex of f and let v_{acc} denote the accept vertex. If u and v are vertices, let $p_{u\to v}$ be the probability that a random walk starting at u reaches v. We use the shorthand $p_{\to v} = p_{v_{\text{start}}\to v}$ and $p_{u\to} = p_{u\to v_{\text{acc}}}$. We use V_i to denote the set of vertices in the i-th layer of the ROBP, where $i \in \{0, 1, \ldots, n\}$.

3 Derandomizing BPL given a hitting set

In this section, we show that the acceptance probability of an arbitrary polynomial width ROBP can be approximated within a small bias in small space, given a certain hitting set. Theorem 1.3 will follow from this.

Theorem 3.1. *Assume there is a $\frac{1}{2}$-hitting set H for width-w' length-n' ROBPs that can be computed in space s. Then the acceptance probability of a given width-w length-n ROBP f can be approximated within a bias $\pm\varepsilon$, in space $O(s+\log\frac{wn}{\varepsilon})$.*

Here $w' = \left\lceil 9\frac{w^3n^2\log(wn)}{\varepsilon^2}\right\rceil, n' = \left\lceil 5\frac{w^3n^4\log(wn)}{\varepsilon^2}\right\rceil$.

Strictly speaking, Theorem 3.1 ought to be phrased in terms of *families* of ROBPs, to make the space bounds meaningful. That is, we assume there is an algorithm that constructs a $\frac{1}{2}$-hitting set for width-w length-n ROBPs, given w and n as inputs, running in space $s(w,n)$. Then given inputs f,ε, Theorem 4.1 should be understood to say that we can estimate $\mathbb{E}[f]$ to within $\pm\varepsilon$ in space $O(s(w',n')+\log(wn/\varepsilon))$.

We are most interested in the case that ε is a small constant, but we remark that when ε is very small, the parameters of Theorem 3.1 could be improved by applying the recent amplification technique by Ahmadinejad et al. [AKM+19].

We first give the derandomization and then give the analysis.

3.1 Derandomization based on a local consistency test

For $x\in\{0,1\}^{n'}$, we interpret it as a concatenation of wn segments. For each $i\in[n]$ and each $v\in V_i$, there is a segment corresponding to v consisting of a concatenation of t sample strings of length i, where t is a power of two satisfying $t\geq 4(\frac{wn}{\varepsilon})^2\log(wn)$. Let $\widehat{p}_{\to v}(x)$ be the fraction of strings that lead to v from the start vertex, among these t sample strings for v. When x is clear, we simply denote it as $\widehat{p}_{\to v}$. Also, for $v\in V_0$, we let $\widehat{p}_{\to v}=1$ if $v=v_{\text{start}}$ and $\widehat{p}_{\to v}=0$ otherwise.

The derandomization conducts a local consistency test $\mathsf{Test}:\{0,1\}^{n'}\to\{0,1\}$ for every $x\in H$ as follows. For all $i\in[n]$, for all $v\in V_i$, check if

$$\left|\widehat{p}_{\to v}-\left(\sum_{u\in V_{i-1}}\widehat{p}_{\to u}\cdot p_{u\to v}\right)\right|\leq\left(1+\sum_{u\in V_{i-1}}p_{u\to v}\right)\varepsilon', \tag{1}$$

where $\varepsilon'=\frac{\varepsilon}{2wn}$. If x passes the checks for all v, then $\mathsf{Test}(x)=1$, otherwise it is 0.

Finally we find an $x\in H$ that passes Test, and output $\widehat{p}_{\to v_{\text{acc}}}(x)$ as the approximation of $\mathbb{E}[f]$.

3.2 Analysis

We now define the "sample verification" function f' of f. For each $x\in\{0,1\}^{n'}$, we set $f'(x)=1$ if and only if for every vertex v in f,

$$|\widehat{p}_{\to v}-p_{\to v}|\leq\varepsilon'. \tag{2}$$

Lemma 3.2. *f' can be computed by a width-w' length-n' ROBP.*

Proof. For each vertex v of f, we construct an ROBP f'_v which simulates f on each sample string and counts how many lead to v. It stores a state of f and a counter value, for a total width of $w\cdot(t+1)$ and a total length $i\cdot t$. f'_v accepts if and only if the counter value is in $[p_{\to v}t-\frac{\varepsilon}{2wn}t, p_{\to v}t+\frac{\varepsilon}{2wn}t]$.

To construct f', we take the conjunction of f'_v, over all v in f. Note that this is a conjunction of ROBPs over disjoint variables. So we can easily see that f' can be computed by an ROBP with width at most $w(t+1)+1\leq\left\lceil 9\frac{w^3n^2\log(wn)}{\varepsilon^2}\right\rceil$, length at most $tw\sum_{i=1}^{n}i\leq\left\lceil 5\frac{w^3n^4\log(wn)}{\varepsilon^2}\right\rceil$. □

Lemma 3.3. *The acceptance probability of f' is at least $\frac{1}{2}$.*

Proof. By the construction of f', for each v of f, there are t uniform random samples. For each sample string, the probability that it leads to v from v_{start} in f is $p_{\to v}$. Hence the expected number of samples leading to v from v_{start} is $p_{\to v}t$. So by Hoeffding's inequality, $\Pr[|\widehat{p}_{\to v}t - p_{\to v}t| \geq \frac{\varepsilon}{2wn}t] \leq 2 \cdot 2^{-2\log(wn)} \leq \frac{2}{(wn)^2}$. There are wn vertices that need to be tested in f. (For $v \in V_0$, the estimate $\widehat{p}_{\to v}$ is always exactly correct.) Thus by a union bound,

$$\Pr\left[\forall v, |\widehat{p}_{\to v} - p_{\to v}| \leq \frac{\varepsilon}{2wn}\right] \geq 1 - \frac{2}{wn}.$$

This is at least $\frac{1}{2}$ when considering n to be at least some large enough constant. So by the definition of f', its acceptance probability is at least $\frac{1}{2}$. □

Lemma 3.4. *For every $x \in \{0,1\}^{n'}$, if $f'(x) = 1$ then $\mathsf{Test}(x) = 1$.*

Proof. For every $i \in [n]$, every $v \in V_i$,

$$p_{\to v} = \sum_{u \in V_{i-1}} p_{\to u} p_{u \to v}, \tag{3}$$

by the structure of ROBP. So

$$\begin{aligned}
\left|\widehat{p}_{\to v} - \sum_{u \in V_{i-1}} \widehat{p}_{\to u} p_{u \to v}\right| &= \left|\widehat{p}_{\to v} - p_{\to v} + p_{\to v} - \sum_{u \in V_{i-1}} \widehat{p}_{\to u} p_{u \to v}\right| \\
&= \left|\widehat{p}_{\to v} - p_{\to v} + \sum_{u \in V_{i-1}} p_{\to u} p_{u \to v} - \sum_{u \in V_{i-1}} \widehat{p}_{\to u} p_{u \to v}\right| && \text{(Equation (3))} \\
&\leq |\widehat{p}_{\to v} - p_{\to v}| + \sum_{u \in V_{i-1}} |p_{\to u} - \widehat{p}_{\to u}|\, p_{u \to v} && \text{(Triangle Inequality)} \\
&\leq \left(1 + \sum_{u \in V_{i-1}} p_{u \to v}\right) \varepsilon'. && \text{(Equation (2))}
\end{aligned}$$

□

Lemma 3.5. *For every $x \in \{0,1\}^{n'}$, if $\mathsf{Test}(x) = 1$ then $|\widehat{p}_{\to v_{\text{acc}}} - p_{\to v_{\text{acc}}}| \leq \varepsilon$.*

Proof. We use induction to show that for the i-th layer of f,

$$\sum_{v \in V_i} |\widehat{p}_{\to v} - p_{\to v}| \leq 2wi\varepsilon'.$$

For the base case, when $i = 0$, it's trivially true since we set $\widehat{p}_{\to v} = p_{\to v}$ for each $v \in V_0$. For the

induction case, assume the hypothesis is true for layer i. Consider layer $i+1$.

$$
\begin{aligned}
&\sum_{v\in V_{i+1}} |\widehat{p}_{\to v} - p_{\to v}| \\
=&\sum_{v\in V_{i+1}} \left|\widehat{p}_{\to v} - \sum_{u\in V_i} p_{\to u}p_{u\to v}\right| && \text{(Equation (3))} \\
=&\sum_{v\in V_{i+1}} \left|\widehat{p}_{\to v} - \sum_{u\in V_i} \widehat{p}_{\to u}p_{u\to v} + \sum_{u\in V_i} \widehat{p}_{\to u}p_{u\to v} - \sum_{u\in V_i} p_{\to u}p_{u\to v}\right| \\
\le&\sum_{v\in V_{i+1}} \left(\left|\widehat{p}_{\to v} - \sum_{u\in V_i} \widehat{p}_{\to u}p_{u\to v}\right| + \left|\sum_{u\in V_i} \widehat{p}_{\to u}p_{u\to v} - \sum_{u\in V_i} p_{\to u}p_{u\to v}\right|\right) && \text{(Triangle Inequality)} \\
\le&\sum_{v\in V_{i+1}} \left|\widehat{p}_{\to v} - \sum_{u\in V_i} \widehat{p}_{\to u}p_{u\to v}\right| + \sum_{v\in V_{i+1}}\sum_{u\in V_i} p_{u\to v}\,|\widehat{p}_{\to u} - p_{\to u}| && \text{(Triangle Inequality)} \\
\le&\sum_{v\in V_{i+1}} \left(1 + \sum_{u\in V_i} p_{u\to v}\right)\varepsilon' + \sum_{v\in V_{i+1}}\sum_{u\in V_i} p_{u\to v}\,|\widehat{p}_{\to u} - p_{\to u}| && (\mathsf{Test}(x)=1) \\
=&2w\varepsilon' + \sum_{u\in V_i} |\widehat{p}_{\to u} - p_{\to u}| && (4) \\
\le&2w\varepsilon' + 2wi\varepsilon' && \text{(Induction hypothesis)} \\
=&2w\cdot(i+1)\cdot\varepsilon'.
\end{aligned}
$$

Here Equation (4) is due to structures of ROBPs. Note that

$$\sum_{v\in V_{i+1}}\sum_{u\in V_i} p_{u\to v} = \sum_{u\in V_i}\sum_{v\in V_{i+1}} p_{u\to v} = w,$$

since for every pair (u,v), $p_{u\to v}$ appears and only appears once in the summation. Also due to the same reasoning,

$$\sum_{v\in V_{i+1}}\sum_{u\in V_i} p_{u\to v}|\widehat{p}_{\to u} - p_{\to u}| = \sum_{u\in V_i}\sum_{v\in V_{i+1}} p_{u\to v}|\widehat{p}_{\to u} - p_{\to u}| = \sum_{u\in V_i} |\widehat{p}_{\to u} - p_{\to u}|.$$

As a result, for the last layer,

$$|\widehat{p}_{\to v_{\text{acc}}} - p_{\to v_{\text{acc}}}| \le \sum_{v\in V_n} |\widehat{p}_{\to v} - p_{\to v}| \le 2wn\varepsilon' = \varepsilon. \qquad \square$$

Lemma 3.6. *The derandomization is in space $O(s + \log \frac{wn}{\varepsilon})$.*

Proof. Since H is computable in space s, for every $x \in H$ we can output any specified bit of it in space $O(s + \log n')$. So when considering the space for computing $\mathsf{Test}(x)$ and $\widehat{p}_{\to v}(x)$, we can just regard x as an input string and only consider working space.

Given vertex v in f, we first consider the space for computing $\widehat{p}_{\to v}$. By the definition of $\widehat{p}_{\to v}$, we can locate the starting position of the t samples for v, taking space $O(\log \frac{wn}{\varepsilon})$. From there, we read the t samples one by one. For each sample, we run f from v_{start} to the layer of v to test if the

sample leads to v. We use a counter c to record the number of samples leading to v. Then compute $\widehat{p}_{\to v}$ as c/t. Since t is a power of two, we can store this number exactly, with no rounding errors. So this step takes space $O(\log(wn)) + O(\log t) = O(\log \frac{wn}{\varepsilon})$. Thus the whole computation is in space $O(\log \frac{wn}{\varepsilon})$.

Next we consider Test. By the definition of Test, for every $i \in [n]$, for each vertex $v \in V_i$, we only need to compute $\widehat{p}_{\to v}$, $\sum_{u \in V_{i-1}} \widehat{p}_{\to u} \cdot p_{u \to v}$ and then test the inequality (1). This again takes space $O(\log \frac{wn}{\varepsilon})$. Note that computing $\mathsf{Test}(x)$ requires *two-way* access to x.

So the overall space of the derandomization is $O(s + \log \frac{wn}{\varepsilon})$. □

Proof of Theorem 3.1. Given a width-w length-n ROBP f, by Lemma 3.2, the function f' can be computed by a width-w' length-n' ROBP. By Lemma 3.3, the acceptance probability of f' is at least $1/2$. Since H is a $\frac{1}{2}$-hitting set for width-w' length-n' ROBPs, there exists $x \in H$ s.t. $f'(x) = 1$. So by Lemma 3.4, there is an $x \in H$ s.t. $\mathsf{Test}(x) = 1$. Hence we can exhaustively search though H to find an x which passes Test. Further, by Lemma 3.5, for this x, $|\, \widehat{p}_{\to v_{\mathsf{acc}}} - p_{\to v_{\mathsf{acc}}} \,| \le \varepsilon$. This shows the derandomization outputs the desired approximation for $p_{\to v_{\mathsf{acc}}}$.

By Lemma 3.6, the derandomization can be done in space $O(s + \log \frac{wn}{\varepsilon})$. □

Theorem 1.3 is directly implied from Theorem 3.1. The proof is straightforward by applying the well known transformation between logspace computations and ROBPs.

4 Deterministic samplers for constant-width ROBPs

In this section, we will show how to use hitting sets to construct deterministic samplers for constant-width ROBPs, thereby proving Theorem 1.7. Most of the work will go toward establishing the following reduction, which is meaningful even for slightly super-constant width.

Theorem 4.1. *Let $w, n \in \mathbb{N}$ and let $\varepsilon > 0$. Assume there is an $(\frac{\varepsilon}{2n})$-hitting set H for width-$(\binom{w}{2}+1)$ length-n ROBPs computable in space s. Then there is a deterministic ε-sampler for width-w length-n ROBPs that runs in space $O(s + w \log(n/\varepsilon))$.*

Like Theorem 3.1, Theorem 4.1 technically ought to be phrased in terms of families of ROBPs. We should also clarify the model of space-bounded oracle algorithm. We assume that the sampler has write-only access to a "query tape" where it can write down an n-bit query string (the query string does not count against the sampler's space complexity). The sampler can then enter a special "query" state, which returns the result of the query into the algorithm's state and clears the query tape. This simple model was perhaps first studied by Ladner and Lynch [LL76].

4.1 Setting up the reduction

Toward proving Theorem 4.1, we begin by setting up some notation. For any ROBP f and a string $x \in \{0,1\}^{\le n}$, let $v_f(x)$ be the vertex reached when f reads x. Furthermore, define

$$p_f(x) = \mathbb{E}[f(x \circ U_{n-|x|})],$$

i.e., $p_f(x) = p_{v_f(x)\to}$.

Now, let $H \subseteq \{0,1\}^n$ be an ε_H-hitting set for width-$(\binom{w}{2}+1)$ ROBPs. For $i \le n$, let H_i be the set of i-bit prefixes of strings in H, i.e., $H_i = \{x_1 x_2 \dots x_i : x \in H\}$. One can verify that H_i is an ε_H-hitting set for width-$(\binom{w}{2}+1)$ length-i ROBPs.

Let f be the width-w ROBP to which we have oracle access. Let λ denote the empty string. Our goal is to estimate $p_f(\lambda)$. For each $i \le n$, define an equivalence relation $\sim$ on $\{0,1\}^i$ by the rule

$$x \sim y \iff \forall z \in H_{n-i}, f(x \circ z) = f(y \circ z).$$

Lemma 4.2. *If $x \sim y$, then $|p_f(x) - p_f(y)| < \varepsilon_H$.*

Proof. Let $i = |x| = |y|$. Define $g\colon \{0,1\}^{n-i} \to \{0,1\}$ by

$$g(z) = f(x \circ z) \oplus f(y \circ z).$$

The function $g(z)$ can be computed by an ROBP of width $\binom{w}{2} + 1$: we have one state in g for each unordered pair of states in f to run the computations $f(x \circ z), f(y \circ z)$ in parallel, along with one additional $\perp$ state in g to indicate that the two computations converged to the same state. If $|p_f(x) - p_f(y)| \ge \varepsilon_H$, then H_{n-i} hits g, hence $x \not\sim y$. □

Let $[x]$ denote the equivalence class of x, so $[x] \subseteq \{0,1\}^{|x|}$. Our deterministic sampler will be based on numbers $\widetilde{p}_f([x]) \in [0,1]$ for each equivalence class $[x]$. The definition of $\widetilde{p}_f$ will ensure that $\widetilde{p}_f([x]) \approx p_f(x)$ for typical values of x, although there might be some anomalous values of x where $\widetilde{p}_f([x]) \not\approx p_f(x)$.

The definition of $\widetilde{p}_f([x])$ is inductive. For the base case, when $x \in \{0,1\}^n$, define $\widetilde{p}_f([x]) = f(x)$. This is well-defined, because $x \sim y \implies f(x) = f(y)$. For the inductive step, suppose $x \in \{0,1\}^i$ with $i < n$. Define

$$\widetilde{p}_f([x]) = \max_{x' \in H_i \cap [x]} \left(\frac{1}{2}\widetilde{p}_f([x' \circ 0]) + \frac{1}{2}\widetilde{p}_f([x' \circ 1]) \right), \tag{5}$$

with the convention that $\widetilde{p}_f([x]) = 0$ if $H_i \cap [x] = \varnothing$. Our sampler will output[3] $\widetilde{p}_f([\lambda])$. (In Section 4.3, we will explain in more detail how to efficiently compute $\widetilde{p}_f([\lambda])$.)

4.2 Correctness

The upper bound on $\widetilde{p}_f([x])$ is straightforward:

Claim 4.3. *For every i, for every $x \in \{0,1\}^{n-i}$,*

$$\widetilde{p}_f([x]) \le p_f(x) + i\varepsilon_H.$$

Proof. We proceed by induction on i. In the base case $i = 0$, $\widetilde{p}_f([x]) = f(x) = p_f(x)$. For the inductive step $i > 0$, we consider two cases. If $H_i \cap [x] = \varnothing$, then $\widetilde{p}_f([x]) = 0$ and the claim is trivial. Otherwise, there is some $x' \in H_i \cap [x]$ such that

$$\begin{aligned}
\widetilde{p}_f([x]) &= \frac{1}{2}\widetilde{p}_f([x' \circ 0]) + \frac{1}{2}\widetilde{p}_f([x' \circ 1]) && \text{(Equation (5))} \\
&\le \frac{1}{2}p_f(x' \circ 0) + \frac{1}{2}p_f(x' \circ 1) + (i-1)\varepsilon_H && \text{(Induction)} \\
&= p_f(x') + (i-1)\varepsilon_H \\
&< p_f(x) + i\varepsilon_H && \text{(Lemma 4.2.)}
\end{aligned}$$

□

[3] Actually the sampler's output differs slightly from $\widetilde{p}_f([\lambda])$ due to rounding errors.

The lower bound is a little more subtle. If u is a vertex in layer i of f, we say that u is *H-reachable* if there is some $x \in H_i$ with $v_f(x) = u$. Otherwise, we say that u is *H-unreachable*. Let $\widetilde{f}$ be a width-$(w+1)$ ROBP obtained from f by replacing all H-unreachable nodes with reject nodes.[4]

Claim 4.4. *For every i, for every $x \in \{0,1\}^{n-i}$,*

$$\widetilde{p}_f([x]) \geq p_{\widetilde{f}}(x).$$

Proof. We proceed by induction on i. In the base case $i = 0$, $\widetilde{p}_f([x]) = f(x) \geq \widetilde{f}(x) = p_{\widetilde{f}}(x)$. For the inductive step $i > 0$, we consider two cases. If f visits some H-unreachable node when it reads x, then $p_{\widetilde{f}}(x) = 0$ and the claim is trivial. Therefore, assume that when f reads x, every node visited is H-reachable. Then there is some $x' \in H_{n-i}$ such that $v_f(x) = v_f(x')$. Of course when f reads x', every node visited is H-reachable, so

$$v_{\widetilde{f}}(x') = v_f(x') = v_f(x) = v_{\widetilde{f}}(x).$$

Therefore,

$$\begin{aligned}
p_{\widetilde{f}}(x) &= p_{\widetilde{f}}(x') \\
&= \frac{1}{2}p_{\widetilde{f}}(x' \circ 0) + \frac{1}{2}p_{\widetilde{f}}(x' \circ 1) \\
&\leq \frac{1}{2}\widetilde{p}_f([x' \circ 0]) + \frac{1}{2}\widetilde{p}_f([x' \circ 1]) && \text{(Induction)} \\
&\leq \widetilde{p}_f(x) && \text{(Equation (5)).}
\end{aligned}$$

(The last inequality uses the fact that $v_f(x) = v_f(x')$ and hence $x \sim x'$.) □

Corollary 4.5. $|\widetilde{p}_f([\lambda]) - \mathbb{E}[f]| \leq n \cdot \varepsilon_H$.

Proof. By Claim 4.3,

$$\widetilde{p}_f([\lambda]) \leq p_f(\lambda) + n \cdot \varepsilon_H = \mathbb{E}[f] + n \cdot \varepsilon_H.$$

In the other direction, by Claim 4.4,

$$\widetilde{p}_f([\lambda]) \geq p_{\widetilde{f}}(\lambda) = \mathbb{E}[\widetilde{f}].$$

Define $g\colon \{0,1\}^n \to \{0,1\}$ by

$$g(x) = 1 \iff \text{when } f \text{ reads } x, \text{ an } H\text{-unreachable node is visited.}$$

Then g can be computed by a width-$(w+1)$ ROBP by a construction very similar to that of $\widetilde{f}$. By construction, g rejects every string in H. Therefore, $\mathbb{E}[g] < \varepsilon_H$. Furthermore, $g(x) = 0 \implies f(x) = \widetilde{f}(x)$. Therefore, $|\mathbb{E}[\widetilde{f}] - \mathbb{E}[f]| < \varepsilon_H$, so $\widetilde{p}_f([\lambda]) > \mathbb{E}[f] - \varepsilon_H$. □

[4] More precisely, add an extra node to each layer labeled $\perp$. In layers prior to the final layer, both outgoing edges from $\perp$ lead to $\perp$, and both outgoing edges from H-unreachable nodes lead to $\perp$. In the final layer, H-unreachable nodes and $\perp$ are reject nodes.

4.3 Efficiently computing $\widetilde{p}_f([\lambda])$

To complete the proof of Theorem 4.1, we just need to show how to efficiently compute $\widetilde{p}_f([\lambda])$. This is fairly straightforward from the definitions; the details follow.

Proof of Theorem 4.1. Say a string $x \in H_i$ is a *representative* if it is the lexicographically first element of $[x] \cap H_i$. Let $x^{(i,1)}, x^{(i,2)}, \ldots$ be an enumeration of the representatives in H_i in lexicographic order. Given i, j, and oracle access to f, one can compute $x^{(i,j)}$ in space $O(s)$.

Our sampler works its way backward through the branching program, starting at layer n and ending with layer 0. The sampler stores data about layer i and uses it when processing layer $i-1$. Specifically, the data stored regarding layer i consists of a list of numbers $p_{i,1}, p_{i,2}, \ldots$, with the interpretation $p_{i,j} = \widetilde{p}_f([x^{(i,j)}])$, or rather $p_{i,j} \approx \widetilde{p}_f([x^{(i,j)}])$ due to rounding error.

For layer n, we can compute this value exactly by setting $p_{i,j} = f(x^{(i,j)})$. Given these values for layer $i+1$, we compute $p_{i,j}$ by the rule

$$p_{i,j} := \max_{\substack{x' \in H_i \cap [x^{(i,j)}] \\ x^{(i+1,j_0)} \sim x' \circ 0 \\ x^{(i+1,j_1)} \sim x' \circ 1}} \left(\frac{1}{2} p_{i+1,j_0} + \frac{1}{2} p_{i+1,j_1} \right), \tag{6}$$

with the convention $p_{i,j} = 0$ if there is no suitable triple (x', j_0, j_1).

The sampler performs the arithmetic in Equation (6) to within $\lceil \log(2n/\varepsilon) \rceil$ bits of precision. This ensures that the rounding error is not too large in each step; by induction, $|p_{i,j} - \widetilde{p}_f([x^{(i,j)}])| \le \frac{\varepsilon(n-i)}{2n}$. The sampler outputs $p_{0,1}$, which is within ε of $\mathbb{E}[f]$ by Corollary 4.5, since $\varepsilon_H = \frac{\varepsilon}{2n}$.

The number of vertices in each layer of f is at most w, so the number of equivalence classes in $\{0,1\}^i$ is also at most w. Therefore, there are at most w representatives in H_i, and hence there are only w numbers $p_{i,j}$ being stored for each layer. Storing those numbers for the layer currently being processed and the layer most recently processed takes $O(w \log(n/\varepsilon))$ bits of space, so overall, the space complexity of the sampler is $O(s + w \log(n/\varepsilon))$ as claimed. □

Interestingly, the sampler of Theorem 4.1 can be implemented to be non-adaptive, because it only queries f at strings of the form $x \circ y$ or $x \circ b \circ z$, where $x \in H_i$, $y \in H_{n-i}$, $b \in \{0,1\}$, and $z \in H_{n-i-1}$.

4.4 Applying the reduction

Proof of Theorem 1.8. Hoza and Zuckerman constructed an $\left(\frac{\varepsilon}{2n}\right)$-hitting set H even for *polynomial-width* ROBPs that can be computed in space $O(\log^2 n + \log(1/\varepsilon))$ [HZ18]. Combining this result with Theorem 4.1 immediately proves Theorem 1.8. □

To prove Theorem 1.7, we must first amplify the assumed $\frac{1}{2}$-hitting set to get an $\left(\frac{\varepsilon}{2n}\right)$-hitting set. This is straightforward, although we must pay a small penalty in terms of width, length, and cardinality.

Lemma 4.6. *Suppose H is a $\frac{1}{2}$-hitting set for width-$(w+1)$ length-(nm) ROBPs. Divide each string $x \in H$ into blocks of length n, $x = x^{(1)} \circ x^{(2)} \circ \cdots \circ x^{(m)}$. Let $H' = \{x^{(i)} : x \in H, i \in [m]\}$. Then H' is a $(\frac{1}{m})$-hitting set for width-w length-n ROBPs.*

Proof. Let f be a width-w length-n ROBP with $\mathbb{E}[f] \ge 1/m$. Define $g \colon (\{0,1\}^n)^m \to \{0,1\}$ by

$$g(x^{(1)} \circ \cdots \circ x^{(m)}) = \bigvee_{i \in [m]} f(x^{(i)}).$$

Then g can be computed by a width-$(w+1)$ ROBP. Furthermore,

$$\mathbb{E}[g] = 1 - (1 - \mathbb{E}[f])^m \geq 1 - \left(1 - \frac{1}{m}\right)^m > \frac{1}{2}.$$

Therefore, H hits g, hence H' hits f. □

Proof of Theorem 1.7. Combine Lemma 4.6 with Theorem 4.1. □

5 Negative result: A barrier for upgrading hitting sets to PRGs

To directly compare hitting sets and PRGs, it is convenient to address the strings in the hitting set using a *hitting set generator* (HSG).

Definition 5.1. *An* ε-HSG *for* $\mathcal{F}$ *is a function* $G\colon \{0,1\}^s \to \{0,1\}^n$ *such that* $G(\{0,1\}^s)$ *is an* ε*-hitting set for* $\mathcal{F}$*.*

In our theorem statements so far, we have been somewhat informal with the distinction between an individual generator vs. a family of generators. Since our negative result is more "meta" than our other results, we will make a precise definition for clarity's sake.

Definition 5.2. *Let* $s(n,t,\varepsilon)$ *be a space-constructible*[5] *function. An* explicit PRG (HSG) family *for width-*w *large-alphabet ROBPs with seed length* s *is a uniform algorithm* G *that takes as input the parameters* n,t,ε *and a string* $y \in \{0,1\}^{s(n,t,\varepsilon)}$ *and outputs a string* $G_{n,t,\varepsilon}(y) \in \{0,1\}^{tn}$*. The algorithm runs in space* $O(s(n,t,\varepsilon))$*, and for each fixed* n,t,ε*, we require that* $G_{n,t,\varepsilon}$ *is an* ε*-PRG (*ε*-HSG) for width-*w *length-*n *ROBPs over the alphabet* $\{0,1\}^t$*.*

The assumption of Theorem 5.3 says that hitting sets can be upgraded into PRGs with essentially no loss: the width parameter remains the same, and the seed length only increases by a constant factor, for any arbitrary setting of n,t,ε. This is only for simplicity's sake. The proof would still go through even if the parameters deteriorated a little when moving from hitting sets to PRGs.

Theorem 5.3. *Let* w *be a constant. Assume that for every* $s(n,t,\varepsilon)$*, if there exists an explicit HSG family for width-*w *large-alphabet ROBPs with seed length* s*, then there exists an explicit PRG family for width-*w *large-alphabet ROBPs with seed length* $O(s)$*. Then for every constant* $\alpha > 0$*, there exists an explicit PRG family for width-*w *ROBPs with seed length*

$$O(t + \log(n/\varepsilon)\log^\alpha n).$$

5.1 From PRGs with moderate error to HSGs with tiny threshold

As outlined in Section 1.5.3, the proof of Theorem 5.3 is based on two reductions. For the first reduction, we show how to convert any PRG with inverse polynomial error into an ε-HSG for any ε. In the regime $n \geq w$, our reduction is a generalization of Hoza and Zuckerman's reduction [HZ18] to the large-alphabet case $t \gg 1$.

Theorem 5.4. *Let* $w,n,t \in \mathbb{N}$ *and let* $\varepsilon > 0$*. Assume there is a* $(\frac{1}{2w^3n^2})$*-PRG* G *for width-*w *length-*n *ROBPs over the alphabet* $\{0,1\}^t$*, with seed length and space complexity bounded by* s*. Then there is an* ε*-hitting set* H *for width-*w *length-*n *ROBPs over the alphabet* $\{0,1\}^t$*, computable in space* $O(s + t + \log(wn/\varepsilon))$*.*

(Just like Theorems 3.1 and 4.1, Theorem 5.4 technically ought to be phrased in terms of families of ROBPs.)

[5] I.e., given n,t,ε, the value $s(n,t,\varepsilon)$ can be computed in space $O(s(n,t,\varepsilon))$.

5.1.1 Construction of the hitting set H

Our hitting set H relies on a hitting set H_{rect} for *combinatorial rectangles* [LLSZ97]. Recall that a combinatorial rectangle over alphabet Γ of dimension r is a function $g\colon \Gamma^r \to \{0,1\}$ of the form $g(x_1,\ldots,x_r) = g_1(x_1) \wedge \cdots \wedge g_r(x_r)$. Without loss of generality, assume $\varepsilon < \frac{1}{w^2n^2}$ and $s \geq t$. The algorithm to enumerate H is as follows.

1. For all $r \in \left\{1,2,\ldots,\left\lfloor \frac{\log(1/\varepsilon)}{\log(wn)} \right\rfloor\right\}$:
   (a) Let $H_{\text{rect}} \subseteq (\{0,1\}^s)^{2r-1}$ be an ε^4-hitting set for combinatorial rectangles over alphabet $\{0,1\}^s$ of dimension $2r-1$.
   (b) For all sequences $(x_1,y_1,x_2,y_2,\ldots,x_{r-1},y_{r-1},x_r) \in H_{\text{rect}}$ and for all sequences of non-negative integers $(n_1,\ldots,n_r)$ satisfying $n_1+n_2+\cdots+n_r = n-r$, output the (nt)-bit string

$$(G(x_1)|_{n_1 t}) \circ (y_1|_t) \circ (G(x_2)|_{n_2 t}) \circ (y_2|_t) \circ \cdots \circ (y_{r-1}|_t) \circ (G(x_r)|_{n_r t}). \tag{7}$$

In Equation (7), the notation $y|_t$ denotes the t-bit prefix of the bitstring y. The key difference between our construction and Hoza and Zuckerman's original hitting set construction [HZ18] is the presence of the strings y_i, which do not pass through the PRG G.

5.1.2 Proof of correctness

Hoza and Zuckerman's reduction was based on a simple structural lemma for ROBPs [HZ18, Lemma 1]. Toward proving the correctness of H, we will now prove a new variant of that lemma, applicable to ROBPs over a large alphabet. For two vertices u, v in an ROBP f, write $u \rightsquigarrow v$ if there is an edge from u to v. Let $\rightsquigarrow^*$ be the reflexive transitive closure of $\rightsquigarrow$, i.e., $u \rightsquigarrow^* v$ if $u = v$ or there is a path from u to v.

The way to think about Lemma 5.5 is to suppose that one is choosing a route from u to v_{acc}. Lemma 5.5 suggests two vertices $v \rightsquigarrow u'$ that one could visit on the way. Item 2 says that it is not difficult to find v. Item 3 says that if one can make it to u', it will be quite a bit easier to find v_{acc} from there. Item 4 says that overall, visiting v and u' is only a mild detour.

In general, in any ROBP over the alphabet Σ, if $v \rightsquigarrow u'$, then $p_{v\to u'} \geq 1/|\Sigma|$. In Hoza and Zuckerman's lemma [HZ18, Lemma 1], they assume $\Sigma = \{0,1\}$, and they use the fact that therefore $p_{v\to u'} \geq \Omega(1)$. At a high level, the reason we need a new structural lemma is that if Σ is large, $p_{v\to u'}$ might be small. Indeed, observe that Lemma 5.5 does not guarantee any lower bound on $p_{v\to u'}$.

Lemma 5.5. *Let f be a width-w, length-n ROBP over any alphabet. Let u be a vertex in f, and assume $0 < p_{u\to} \leq \frac{1}{wn}$. Then there is a pair of vertices (v,u') in f such that:*

1. $u \rightsquigarrow^* v \rightsquigarrow u'$.
2. $p_{u\to v} \geq \frac{1}{w^3n^2}$.
3. $p_{u'\to} \geq wn \cdot p_{u\to}$.
4. $p_{u\to v} \cdot p_{v\to u'} \cdot p_{u'\to} \geq \frac{p_{u\to}}{w^2 n}$.

Proof. Suppose some pair (v,u') satisfies Item 1, but it violates Item 4. For such a pair, if we take a random walk from u, the probability that we visit v, u', and v_{acc} is less than $\frac{p_{u\to}}{w^2n}$. The number of such pairs is at most w^2n, so by the union bound, when we start at u and read random bits, the

probability that we visit *any* such pair and v_{acc} is less than $p_{u\to}$. Therefore, there is some path from u to v_{acc} that never visits such a pair.

Let u' be the first vertex along that path that satisfies Item 3. (Such a u' exists, because if nothing else we can let $u' = v_{\text{acc}}$.) Let v be the vertex immediately preceding u' in the path. (This makes sense, because $p_{u\to} < wn \cdot p_{u\to}$, so $u' \neq u$.) This pair clearly satisfies Items 1, 3 and 4; all that remains is to verify Item 2. Indeed,

$$\begin{aligned}\frac{p_{u\to}}{p_{u\to v}} &\le w^2 n \cdot p_{v\to u'} \cdot p_{u'\to} && \text{(Item 4)}\\ &\le w^2 n \cdot p_{v\to}\\ &< w^2 n \cdot wn \cdot p_{u\to},\end{aligned}$$

where the last inequality holds because u' is the *first* vertex in the path satisfying Item 3, and v precedes u', so v must not satisfy Item 3. Rearranging completes the proof. □

Corollary 5.6. *Let $0 < \varepsilon \le \frac{1}{wn}$. Let f be a width-w, length-n ROBP over any alphabet with $\mathbb{E}[f] \ge \varepsilon$. Then there is a sequence of vertices*

$$v_{\text{start}} = u_1 \rightsquigarrow^* v_1 \rightsquigarrow u_2 \rightsquigarrow^* v_2 \rightsquigarrow \cdots \rightsquigarrow u_r \rightsquigarrow^* v_r = v_{\text{acc}}$$

such that:

1. *For every i, $p_{u_i\to v_i} \ge \frac{1}{w^3 n^2}$.*
2. *$r \le \frac{\log(1/\varepsilon)}{\log(wn)}$.*
3. *$p_{u_1\to v_1} \cdot p_{v_1\to u_2} \cdot p_{u_2\to v_2} \cdots p_{v_{r-1}\to u_r} \cdot p_{u_r\to v_r} \ge \varepsilon^3$.*

Proof. We define the sequence inductively, starting with $u_1 = v_{\text{start}}$. Assume we've defined $u_1, v_1, u_2, v_2, \ldots, u_i$. If $p_{u_i\to} \ge \frac{1}{w^3 n^2}$, then set $r = i$, set $v_i = v_{\text{acc}}$, and terminate the sequence. Otherwise, let (v_i, u_{i+1}) be the vertices provided by plugging $u = u_i$ into Lemma 5.5.

Item 1 of Lemma 5.5 implies that $u_i \rightsquigarrow^* v_i$ and $v_i \rightsquigarrow u_{i+1}$. Item 1 is guaranteed by Item 2 of Lemma 5.5 and the termination condition. By Item 3 of Lemma 5.5, $p_{u_{i+1}\to} \ge wn \cdot p_{u_i\to}$, which implies Item 2. Finally, iteratively applying Item 4 of Lemma 5.5 shows that

$$p_{u_1\to v_1} \cdot p_{v_1\to u_2} \cdot p_{u_2\to v_2} \cdots p_{v_{r-1}\to u_r} \cdot p_{u_r\to v_r} \ge \frac{p_{u_1\to}}{(w^2 n)^r} \ge \varepsilon^3,$$

i.e., Item 3 holds. □

We are now ready to complete the proof of correctness of our hitting set H.

Claim 5.7. *If f is a width-w length-n ROBP over the alphabet $\{0,1\}^t$ with $\mathbb{E}[f] \ge \varepsilon$, then $f^{-1}(1) \cap H \neq \varnothing$.*

Proof. Let $u_1 \rightsquigarrow^* v_1 \rightsquigarrow \cdots \rightsquigarrow u_r \rightsquigarrow^* v_r$ be the sequence of vertices guaranteed by Corollary 5.6. Let n_i be the distance from u_i to v_i. Let $g\colon (\{0,1\}^s)^{2r-1} \to \{0,1\}$ be the following combinatorial rectangle:

$$\begin{aligned}&g(x_1, y_1, x_2, y_2, \ldots, x_{r-1}, y_{r-1}, x_r) = 1 \iff\\ &\qquad \forall i \in [r],\ G(x_i)|_{n_i t} \text{ leads from } u_i \text{ to } v_i \ \text{ and } \ \forall i \in [r-1],\ y_i|_t \text{ leads from } v_i \text{ to } u_{i+1}.\end{aligned}$$

By Item 1 of Corollary 5.6, $p_{u_i \to v_i} \geq \frac{1}{w^3 n^2}$. Since G has error $\frac{1}{2w^3n^2}$,

$$\Pr[G(U) \text{ leads from } u_i \text{ to } v_i] \geq \frac{1}{2} p_{u_i \to v_i}.$$

Therefore, by Item 3 of Corollary 5.6, $\mathbb{E}[g] \geq \varepsilon^3 \cdot 2^{-r} \geq \varepsilon^4$. Therefore, there is some sequence $(x_1, y_1, \ldots, y_{r-1}, x_r) \in H_{\text{rect}}$ that hits g. By construction, the corresponding element of H is accepted by f. □

5.1.3 Efficiency

Proof of Theorem 5.4. To complete the proof of Theorem 5.4, let us analyze the space complexity of H. The number r can be stored using $O(\log\log(1/\varepsilon))$ bits of space. Using a construction by Linial, Luby, Saks, and Zuckerman [LLSZ97], because of our chosen value of r, we can enumerate H_{rect} in space $O(s + \log(1/\varepsilon))$. The integers $n_1, \ldots, n_r$ can be straightforwardly stored using $O(r \log n) = O(\log(1/\varepsilon))$ bits of space. Thus, overall, the space complexity is $O(s + \log(1/\varepsilon))$. (Recall that we assumed without loss of generality that $\varepsilon < \frac{1}{w^2n^2}$ and $s \geq t$.) □

5.2 Application: Unconditional improved hitting sets for large-alphabet ROBPs

As outlined in Section 1.5.3, plugging the class INW generator [INW94] into the reduction of Theorem 5.4 already gives something interesting: an improved hitting set for large-alphabet ROBPs, even of polynomial width.

Corollary 5.8. *Let $w, n, t \in \mathbb{N}$ and let $\varepsilon > 0$. There is an ε-hitting set H for width-w length-n ROBPs over the alphabet $\{0,1\}^t$, computable in space $O(t + \log(wn)\log n + \log(1/\varepsilon))$.*

5.3 Trading a good dependence on ε for a good dependence on n

Recall that to prove Theorem 5.3, we must (conditionally) construct a PRG with a good dependence on n. So far, unconditionally, Theorem 5.4 has provided us with an HSG with a good dependence on ε. The assumption of Theorem 5.3 allows us to convert that HSG into a PRG with the same seed length, $O(t + \log^2 n + \log(1/\varepsilon))$ (for width w, a constant). In this section, we show how to convert that PRG into another PRG with seed length $O(t + \log^{3/2} n + \log(1/\varepsilon)\sqrt{\log n})$, i.e., we improve the dependence on n at the expense of a worse dependence on ε. That follows from setting $\alpha = 1/2$ in the following more general reduction.

Lemma 5.9. *Let $\alpha \in (0,1)$ be a constant. Let $w, n, t \in \mathbb{N}$ and $\varepsilon > 0$. Define $m = \left\lceil 2^{(\log n)^{1-\alpha}} \right\rceil$ and $d = \lceil C \log(n/\varepsilon) \rceil$, where C is an appropriate constant. Assume there is an $(\frac{\varepsilon}{4n})$-PRG G for width-w length-m ROBPs over the alphabet $\{0,1\}^d$ with seed length and space complexity bounded by s. Then there is an ε-PRG G' for width-w length-n ROBPs over the alphabet $\{0,1\}^t$ with seed length and space complexity $O(t + s \cdot \log^\alpha n + \log(wn/\varepsilon))$.*

As usual, Lemma 5.9 should technically be phrased in terms of families of ROBPs. As suggested in Section 1.5.3, the proof of Lemma 5.9 is not particularly novel. It is an application of traditional seed-recycling techniques, similar to classic constructions of PRGs for space-bounded computation [Nis92, INW94, NZ96]. Our construction and analysis are especially similar to Armoni's work [Arm98].

One difference is that we use randomized samplers rather than extractors for convenience; in this respect, our construction is similar to a variant of the INW generator [INW94] described by Braverman, Cohen, and Garg [BCG18] as a warm-up to their main construction. In particular, we rely on the following randomized sampler by Goldreich and Wigderson [GW97].

Theorem 5.10 ([GW97, Lemma 6.6])**.** *For all $t \in \mathbb{N}$, $\delta > 0$, there exists a function* $\mathsf{Samp}\colon \{0,1\}^t \times \{0,1\}^{O(\log(1/\delta))} \to \{0,1\}^t$ *such that for any*[6] *function* $f\colon \{0,1\}^t \to [0,1]$*,*

$$\Pr_x\left[\left|\mathop{\mathbb{E}}_y[f(\mathsf{Samp}(x,y))] - \mathbb{E}[f]\right| \le \delta\right] \ge 1-\delta.$$

Furthermore, given t, δ, x, y as inputs, $\mathsf{Samp}(x,y)$ can be computed in space $O(t)$.

We will recursively use the following basic PRG, which stretches $t + dn$ bits to tn bits. It might be helpful to think of the case $t = 100d$.

Lemma 5.11. *Let t, δ be arbitrary, and let* $\mathsf{Samp}\colon \{0,1\}^t \times \{0,1\}^d \to \{0,1\}^t$ *be the randomized sampler of Theorem 5.10. Define* $G_0\colon \{0,1\}^t \times (\{0,1\}^d)^n \to (\{0,1\}^t)^n$ *by*

$$G_0(x, z_1 \circ \cdots \circ z_n) = \mathsf{Samp}(x, z_1) \circ \cdots \circ \mathsf{Samp}(x, z_n).$$

Then G_0 fools width-w length-n ROBPs over the alphabet $\{0,1\}^t$ with error $\delta w^2 n$.

The proof of Lemma 5.11 is straightforward, and we omit it. When reading the proof of Lemma 5.9, it might be helpful to keep in mind that all "x" variables are strings of length t, all "y" variables are strings of length s, and all "z" variables are strings of length d.

Proof of Lemma 5.9. Define $n_i = n/m^i$. For simplicity, we ignore rounding issues, i.e., we assume that n_i is an integer and that $m = 2^{(\log n)^{1-\alpha}}$ exactly. Let $\delta = \frac{\varepsilon}{4w^2 n}$, and let $d = O(\log(wn/\varepsilon))$ be the length of the second input to the function Samp of Theorem 5.10. We will recursively define a sequence of PRGs

$$G_i\colon \{0,1\}^t \times (\{0,1\}^s)^i \times (\{0,1\}^d)^{n_i} \to (\{0,1\}^t)^n.$$

The base case $i = 0$ is the basic PRG of Lemma 5.11:

$$G_0(x, z_1 \circ \cdots \circ z_n) = \mathsf{Samp}(x, z_1) \circ \cdots \circ \mathsf{Samp}(x, z_n).$$

For the inductive step $i > 0$, we define[7]

$$G_i(x, y_1 \circ \cdots \circ y_i, z_1 \circ \cdots \circ z_{n_i}) = G_{i-1}(x, y_1 \circ \cdots \circ y_{i-1}, G(\mathsf{Samp}(y_i, z_1)) \circ \cdots \circ G(\mathsf{Samp}(y_i, z_{n_i}))),$$

where G is the given PRG. To analyze these generators, let f be a width-w length-n ROBP over the alphabet $\{0,1\}^t$. For each i and each fixing of $x, y_1, \ldots, y_i$, define

$$g^{(x,y_1,\ldots,y_i)}(z_1 \circ \cdots \circ z_{n_i}) = f(G_i(x, y_1 \circ \cdots \circ y_i, z_1 \circ \cdots \circ z_{n_i}))$$
$$h^{(x,y_1,\ldots,y_i)}(y'_1 \circ \cdots \circ y'_{n_{i+1}}) = f(G_i(x, y_1 \circ \cdots \circ y_i, G(y'_1) \circ \cdots \circ G(y'_{n_{i+1}})).$$

These functions are related to one another by the rules

$$h^{(x,y_1,\ldots,y_i)}(y'_1 \circ \cdots \circ y'_{n_{i+1}}) = g^{(x,y_1,\ldots,y_i)}(G(y'_1) \circ \cdots \circ G(y'_{n_{i+1}})) \tag{8}$$
$$g^{(x,y_1,\ldots,y_i)}(z_1 \circ \cdots \circ z_{n_i}) = h^{x,y_1,\ldots,y_{i-1}}(\mathsf{Samp}(y_i, z_1) \circ \cdots \circ \mathsf{Samp}(y_i, z_{n_i})). \tag{9}$$

[6] Goldreich and Wigderson analyze the case that f is $\{0,1\}$-valued, but the $[0,1]$-valued case automatically follows with only a quadratic loss in δ.

[7] Note that strictly speaking, we are using *two* instantiations of Samp. In the base case, Samp has output length t, whereas in the inductive step, Samp has output length s. Hopefully, using the same name Samp for both samplers will not cause confusion.

This shows by induction on i that each g function can be computed by a width-w ROBP over the alphabet $\{0,1\}^d$ and each h function can be computed by a width-w ROBP over the alphabet $\{0,1\}^s$.

Let us now show by induction on i that G_i fools f with error $(\delta w^2 + \varepsilon_G) \cdot \sum_{j=0}^{i} n_j$, where ε_G is the error of G. The base case $i = 0$ is already established by Lemma 5.11. For the inductive step, we have

$$\begin{aligned}
&\mathop{\mathbb{E}}_{\substack{x\\ y_1,\dots,y_i\\ z_1,\dots,z_{n_i}}} [f(G_i(x, y_1 \circ \cdots \circ y_i, z_1 \circ \cdots \circ z_{n_i}))] \\
&= \mathop{\mathbb{E}}_{\substack{x\\ y_1,\dots,y_i\\ z_1,\dots,z_{n_i}}} [g^{(x,y_1,\dots,y_i)}(z_1 \circ \cdots \circ z_{n_i})] \\
&= \mathop{\mathbb{E}}_{\substack{x\\ y_1,\dots,y_i\\ z_1,\dots,z_{n_i}}} [h^{(x,y_1,\dots,y_{i-1})}(\mathsf{Samp}(y_i, z_1) \circ \cdots \circ \mathsf{Samp}(y_i, z_{n_i}))] && \text{(Equation (9))} \\
&\le \mathop{\mathbb{E}}_{\substack{x\\ y_1,\dots,y_{i-1}\\ y'_1,\dots,y'_{n_i}}} [h^{(x,y_1,\dots,y_{i-1})}(y'_1 \circ \cdots \circ y'_{n_i})] + \delta w^2 n_i && \text{(Lemma 5.11)} \\
&= \mathop{\mathbb{E}}_{\substack{x\\ y_1,\dots,y_{i-1}\\ y'_1,\dots,y'_{n_i}}} [g^{(x,y_1,\dots,y_{i-1})}(G(y'_1) \circ \cdots \circ G(y'_{n_i}))] + \delta w^2 n_i && \text{(Equation (8))} \\
&\le \mathop{\mathbb{E}}_{\substack{x\\ y_1,\dots,y_{i-1}\\ z_1,\dots,z_{n_{i-1}}}} [g^{(x,y_1,\dots,y_{i-1})}(z_1 \circ \cdots \circ z_{n_{i-1}})] + (\delta w^2 + \varepsilon_G) \cdot n_i \\
&= \mathop{\mathbb{E}}_{\substack{x\\ y_1,\dots,y_{i-1}\\ z_1,\dots,z_{n_{i-1}}}} [f(G_{i-1}(x, y_1 \circ \cdots \circ y_{i-1}, z_1 \circ \cdots \circ z_{n_{i-1}}))] + (\delta w^2 + \varepsilon_G) \cdot n_i.
\end{aligned}$$

The lower bound follows the same argument. Let $r = \log^\alpha n$ and $G' = G_r$. Then G' fools f with error

$$(\delta w^2 + \varepsilon_G) \cdot \sum_{i=0}^{r} n_i \le (\delta w^2 + \varepsilon_G) \cdot n \cdot \sum_{i=0}^{\infty} m^{-i} \le 2n \cdot (\delta w^2 + \varepsilon_G) \le \varepsilon.$$

Furthermore, the seed length of G' is $t + rs + d$ as claimed, and the space complexity of G' is clearly also $O(t + rs + d)$. □

5.4 Putting things together to prove Theorem 5.3

Proof of Theorem 5.3. We will show by induction on a that for each constant $a \in \mathbb{N}$, there is an explicit PRG family for width-w ROBPs with seed length $O(t + \log(n/\varepsilon) \log^{1/a} n)$. The base case $a = 1$ holds unconditionally – this is the seed length of the classic INW generator [INW94].

For the inductive step, suppose $a > 1$. Let t, n, ε be arbitrary; we will construct an ε-PRG for width-w length-n ROBPs over the alphabet $\{0,1\}^t$. Define $\alpha = 1/a$. Let $m = 2^{(\log n)^{1-\alpha}}$ and $d = C \log(n/\varepsilon)$, as in Lemma 5.9.

By induction, there is a $(\frac{1}{2w^3m^2})$-PRG G for width-w length-m ROBPs over the alphabet $\{0,1\}^d$, with seed length and space complexity bounded by $O(d + \log^{1+\frac{1}{a-1}} m)$. Now,

$$\log^{1+\frac{1}{a-1}} m = (\log n)^{(1-\frac{1}{a}) \cdot (1+\frac{1}{a-1})} = \log n,$$

so G has seed length and space complexity bounded by $O(\log(n/\varepsilon))$. Plugging G into Theorem 5.4, we get an $(\frac{\varepsilon}{4n})$-hitting set H for width-w length-m ROBPs over the alphabet $\{0,1\}^d$, computable in

space $O(\log(n/\varepsilon))$. Now we use our assumption to convert H into a PRG G' with exactly the same parameters. Finally, plugging G' into Lemma 5.9 gives the desired PRG. □

6 Directions for further research

In this paper, we have shown that hitting sets for **RL** would derandomize **BPL**. Constructing a hitting set is the most natural way to prove $\mathbf{L} = \mathbf{RL}$, but there are also other approaches. In general, does $\mathbf{L} = \mathbf{RL}$ imply $\mathbf{L} = \mathbf{BPL}$? In the polynomial-time setting, the "promise" variant of this question has been answered in the affirmative, i.e., $\mathbf{prP} = \mathbf{prRP} \implies \mathbf{P} = \mathbf{BPP}$ [BF99]. Does $\mathbf{prL} = \mathbf{prRL}$ imply $\mathbf{L} = \mathbf{BPL}$? Or relaxing the challenge even further, does $\mathbf{L} = \mathbf{NL}$ imply $\mathbf{L} = \mathbf{BPL}$?

We gave two different algorithms for estimating the expectation of an ROBP given a hitting set, one suited for $w = \text{poly}(n)$ (Theorem 3.1) and one suited for $w = O(1)$ (Theorem 4.1). What about the case $n = \text{polylog}\, w$? Unconditionally, there are optimal hitting sets known in this regime [AKS87, NZ96, HZ18]. Given such an ROBP f as input, is it possible to compute $\mathbb{E}[f] \pm \frac{1}{w}$ in space $O(\log w)$? An affirmative answer would imply that any space-s decision algorithm that uses n random bits could be simulated by another space-$O(s)$ algorithm using only $O(n/s^c)$ random bits, where c is an arbitrarily large constant.

Recently, Meka, Reingold, and Tal constructed a PRG for width-3 ROBPs with seed length $\widetilde{O}(\log n \log(1/\varepsilon))$ [MRT19]. This is near-optimal when ε is not too small, but for $\varepsilon = 1/n$ it is worse than Nisan's PRG [Nis92]. On the other hand, there is an explicit hitting set for width-3 ROBPs with near-optimal seed length $\widetilde{O}(\log(n/\varepsilon))$ [GMR+12]. Can one construct an explicit deterministic sampler for width-3 ROBPs with near-optimal seed length? Unfortunately, to produce a deterministic sampler for width-3 ROBPs, Theorem 4.1 would require a hitting set for width-4 ROBPs.

Assuming the existence of a log-space hitting set for polynomial-width ROBPs, is it possible to construct a log-space deterministic sampler for polynomial-width ROBPs?

Recall that PRPDs are superior to deterministic samplers (see Figure 1). Is it possible to improve Theorem 1.7 so that it concludes with a PRPD rather than a mere deterministic sampler?

7 Acknowledgments

We thank David Zuckerman for helpful discussions and for comments on a draft of this paper. We thank Dean Doron and Pooya Hatami for valuable early discussions about this research project.

References

[ACR96] Alexander E. Andreev, Andrea E. F. Clementi, and José D. P. Rolim. Hitting sets derandomize BPP. In *Automata, languages and programming (Paderborn, 1996)*, volume 1099 of *Lecture Notes in Comput. Sci.*, pages 357–368. Springer, Berlin, 1996.

[ACRT99] Alexander E. Andreev, Andrea E. F. Clementi, José D. P. Rolim, and Luca Trevisan. Weak random sources, hitting sets, and BPP simulations. *SIAM Journal on Computing*, 28(6):2103–2116, 1999.

[AKM+19] AmirMahdi Ahmadinejad, Jonathan Kelner, Jack Murtagh, John Peebles, Aaron Sidford, and Salil Vadhan. High-precision estimation of random walks in small space. *arXiv preprint arXiv:1912.04524*, 2019.

[AKS87] Miklós Ajtai, János Komlós, and Endre Szemerédi. Deterministic simulation in LOGSPACE. In *Proceedings of the 19th Annual Symposium on Theory of Computing*, pages 132–140. ACM, 1987.

[Arm98] Roy Armoni. On the derandomization of space-bounded computations. In *Proceedings of the 2nd International Workshop on Randomization and Computation*, volume 1518 of *Lecture Notes in Computer Science*, pages 47–59. Springer, Berlin, 1998.

[BCG18] Mark Braverman, Gil Cohen, and Sumegha Garg. Hitting sets with near-optimal error for read-once branching programs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 353–362, New York, NY, USA, 2018. ACM.

[BDVY13] Andrej Bogdanov, Zeev Dvir, Elad Verbin, and Amir Yehudayoff. Pseudorandomness for width-2 branching programs. *Theory of Computing*, 9:283–292, 2013.

[BF99] H. Buhrman and L. Fortnow. One-sided versus two-sided error in probabilistic computation. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, volume 1563, pages 100–109. Berlin, 1999.

[Bog05] Andrej Bogdanov. Pseudorandom generators for low degree polynomials. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 21–30, 2005.

[GMR+12] Parikshit Gopalan, Raghu Meka, Omer Reingold, Luca Trevisan, and Salil Vadhan. Better pseudorandom generators from milder pseudorandom restrictions. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2012)*, pages 120–129. IEEE, 2012.

[GVW11] Oded Goldreich, Salil Vadhan, and Avi Wigderson. Simplified derandomization of BPP using a hitting set generator. In *Studies in complexity and cryptography*, volume 6650 of *Lecture Notes in Comput. Sci.*, pages 59–67. Springer, Heidelberg, 2011.

[GW97] Oded Goldreich and Avi Wigderson. Tiny families of functions with random properties: a quality-size trade-off for hashing. In *Proceedings of the Workshop on Randomized Algorithms and Computation (Berkeley, CA, 1995)*, volume 11, pages 315–343, 1997.

[HU17] William M. Hoza and Chris Umans. Targeted pseudorandom generators, simulation advice generators, and derandomizing logspace. In *Proceedings of the 49th Annual Symposium on Theory of Computing*, pages 629–640. ACM, New York, 2017.

[HZ18] William M. Hoza and David Zuckerman. Simple optimal hitting sets for small-success RL. In *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2018)*. IEEE, 2018.

[INW94] Russell Impagliazzo, Noam Nisan, and Avi Wigderson. Pseudorandomness for network algorithms. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing*, pages 356–364. ACM, 1994.

[IW97] Russell Impagliazzo and Avi Wigderson. P = BPP if E requires exponential circuits: Derandomizing the XOR lemma. In *Proceedings of the Twenty-ninth Annual ACM Symposium on Theory of Computing*, STOC '97, pages 220–229, New York, NY, USA, 1997. ACM.

[LL76] Richard E. Ladner and Nancy A. Lynch. Relativization of questions about log space computability. *Math. Systems Theory*, 10(1):19–32, 1976.

[LLSZ97] Nathan Linial, Michael Luby, Michael Saks, and David Zuckerman. Efficient construction of a small hitting set for combinatorial rectangles in high dimension. *Combinatorica*, 17(2):215–234, 1997.

[MRT19] Raghu Meka, Omer Reingold, and Avishay Tal. Pseudorandom generators for width-3 branching programs. In *STOC'19—Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 626–637. ACM, New York, 2019.

[Nis92] Noam Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992.

[Nis93] Noam Nisan. On read-once vs. multiple access to randomness in logspace. *Theoretical Computer Science*, 107(1):135–144, 1993.

[NZ96] Noam Nisan and David Zuckerman. Randomness is linear in space. *Journal of Computer and System Sciences*, 52(1):43–52, 1996.

[ŠŽ11] Jiří Šíma and Stanislav Žák. Almost k-wise independent sets establish hitting sets for width-3 1-branching programs. In *Computer science—theory and applications*, volume 6651 of *Lecture Notes in Comput. Sci.*, pages 120–133. Springer, Heidelberg, 2011.

A Derandomizing BPP given a hitting set

Theorem A.1 ([ACR96])**.** *Assume that for every $s, n \in \mathbb{N}$, there is a $\frac{1}{2}$-hitting set $H_{s,n}$ for size-s circuits on n input bits that can be computed in time* $\mathrm{poly}(s, n)$*. Then* $\mathbf{P} = \mathbf{BPP}$*.*

Proof. By naïve amplification, we may assume that the randomized algorithm has failure probability 2^{-N}, where N is the *input* length. Let C be a size-n circuit on n input bits describing the action of this algorithm on its random bits, so $n = \mathrm{poly}(N)$ and we are trying to distinguish the cases $\mathbb{E}[C] \leq 2^{-N}$ vs. $\mathbb{E}[C] \geq 1 - 2^{-N}$. Our algorithm accepts if and only if there exists $x \in H_{n^c,n}$ such that for all $y \in H_{3n,n}$, $C(x \oplus y) = 1$. Here, c is a suitable constant that will become clear later. The runtime is clearly $\mathrm{poly}(N)$.

For the correctness proof, first suppose $\mathbb{E}[C] \leq 2^{-N}$. For any fixed x, the function $y \mapsto \neg C(x \oplus y)$ has expectation at least $1 - 2^{-N}$ and can be computed by a circuit of size $3n$. Therefore, there is some $y \in H_{3n,n}$ such that $C(x \oplus y) = 0$, and hence our algorithm rejects. Conversely, suppose $\mathbb{E}[C] \geq 1 - 2^{-N}$. Consider sampling $x \in \{0,1\}^n$ and $y \in H_{3n,n}$ uniformly at random. Since x is uniform, $\mathbb{E}_{x,y}[\neg C(x \oplus y)] \leq 2^{-N}$. By Markov's inequality,

$$\Pr_{x \in \{0,1\}^n}\left[\mathop{\mathbb{E}}_{y \in H_{3n,n}}[\neg C(x \oplus y)] < 2 \cdot 2^{-N}\right] > 1/2.$$

Since $H_{3n,n}$ can be computed in polynomial time, $|H_{3n,n}| \leq \mathrm{poly}(N)$. Therefore, if $\mathbb{E}_y[\neg C(x \oplus y)] < 2 \cdot 2^{-N}$, then in fact $\mathbb{E}_y[\neg C(x \oplus y)] = 0$, provided N is sufficiently large. Therefore,

$$\Pr_{x \in \{0,1\}^n}[\forall y \in H_{3n,n}, C(x \oplus y) = 1] > 1/2.$$

Given input x, the predicate $\forall y \in H_{3n,n}, C(x \oplus y) = 1$ can be computed by a circuit of size n^c for some suitable constant c. Therefore, there is some $x \in H_{n^c,n}$ that hits that circuit. □

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il