



Near-Optimal Derandomization of Medium-Width Branching Programs*

Aaron (Louie) Putterman
aputterman@g.harvard.edu
Harvard University
USA

Edward Pyne
epyne@mit.edu
Massachusetts Institute of Technology
USA

ABSTRACT

We give a deterministic algorithm to estimate the expectation of a read-once branching program of length n and width w in space

$$\tilde{O}\left(\log n + \sqrt{\log n \cdot \log w}\right). \quad (1)$$

In particular, we obtain a nearly optimal space $\tilde{O}(\log n)$ derandomization of programs up to width $w = 2^{\sqrt{\log n}}$. Previously, the best known space complexity for this problem was

$$O\left(\min\left\{\log n \cdot \log w, \log^{3/2} n + \sqrt{\log n \cdot \log w}\right\}\right)$$

via the classic algorithm of Savitch (JCSS 1970) and Saks and Zhou (JCSS 1999), which only achieve space $\tilde{O}(\log n)$ for $w = \text{polylog}(n)$.

We prove this result by showing that a variant of the Saks-Zhou algorithm developed by Cohen, Doron, and Sberlo (ECCC 2022) still works without executing one of the steps in the algorithm, the so-called “random shift step.” This allows us to extend their algorithm from computing the n th power of a $w \times w$ stochastic matrix to multiplying n distinct $w \times w$ stochastic matrices with no degradation in space consumption. In the regime where $w \geq n$, we also show that our approach can achieve parameters matching those of the original Saks-Zhou algorithm (with no $\log\log$ factors). Finally, we show that for $w \leq 2^{\sqrt{\log n}}$, an algorithm even simpler than our algorithm and that of Saks and Zhou achieves space $O(\log^{3/2} n)$.

CCS CONCEPTS

• **Theory of computation** → **Pseudorandomness and derandomization**.

KEYWORDS

Pseudorandomness, Space-Bounded Computation, Branching Programs

ACM Reference Format:

Aaron (Louie) Putterman and Edward Pyne. 2023. Near-Optimal Derandomization of Medium-Width Branching Programs. In *Proceedings of the*

*A.P. was supported under the Simons Investigator Fellowship of Boaz Barak, NSF grant DMS-2134157, DARPA grant W911NF2010021, and DOE grant DE-SC0022199. E.P. was supported by an Akamai Presidential Fellowship.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '23, June 20–23, 2023, Orlando, FL, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9913-5/23/06...\$15.00

<https://doi.org/10.1145/3564246.3585108>

55th Annual ACM Symposium on Theory of Computing (STOC '23), June 20–23, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3564246.3585108>

1 INTRODUCTION

There have been over four decades of work towards derandomizing space-bounded computation, i.e. proving $\text{BPL} = \text{L}$. A central, extensively studied problem that is “BPL-complete” is to estimate the expectation of a read-once branching program of length n and width $w = n$. A long line of research has attacked $\text{BPL} = \text{L}$ via pseudorandom generators and other black-box derandomization techniques [2, 3, 14, 15, 19, 20].

Here, however, our focus is on the *non-black-box* setting, where the problem is equivalent to computing an approximate product of n stochastic matrices $M_1, \dots, M_n \in \mathbb{R}^{w \times w}$, where M_i corresponds to the transition probabilities of the branching program from layer $i - 1$ to layer i . From the classical work of Savitch [25] it is easy to derive a space $O(\log^2 n)$ algorithm for this problem (where $w = n$) via a divide and conquer approach.

In 1995, Saks and Zhou gave a breakthrough algorithm that achieved space $O(\log^{3/2} n)$ for the same problem:

THEOREM 1.1 ([24]). *There is a deterministic algorithm such that given stochastic matrices $M_1, \dots, M_n \in \mathbb{R}^{w \times w}$, the algorithm computes a $1/n$ entrywise approximation to $M_1 \cdots M_n$. The algorithm runs in space*

$$O\left(\log^{3/2} n + \sqrt{\log n \cdot \log w}\right).$$

In 2022, Cohen, Doron, and Sberlo gave an algorithm that improved on Saks and Zhou in the case where $w \ll n$, and we add the additional constraint that all the stochastic matrices are equal:

THEOREM 1.2 ([10]). *There is a deterministic algorithm such that given a stochastic matrix $M \in \mathbb{R}^{w \times w}$, the algorithm computes a $1/n$ entrywise approximation to M^n . The algorithm runs in space*

$$\tilde{O}\left(\log n + \sqrt{\log n \cdot \log w}\right).$$

Unfortunately, Cohen et al.’s result does not seem to improve on Saks and Zhou for computing the product of n distinct $w \times w$ stochastic matrices, and thus does not improve on the space complexity of derandomizing branching programs. We remark that for branching programs of width $w = o(n)$, the restriction that all transition functions are the same is severe: there are simple functions computable by ordered branching programs of width 2 that cannot be computed by “single transition” ordered programs of width $n/2 - 1$ (see Appendix B). Thus, for estimating the expectation

¹We assume all input matrices have $O(\log n)$ bit entries for simplicity.

of a width w , length n read-once branching program, the best space complexity remains

$$\min \left\{ O(w \log n), \quad O(\log n \cdot \log w) + \tilde{O}(\log n), \right. \\ \left. O\left(\log^{1/2-o(1)}(n)(\log n + \log w)\right) \right\}$$

, where the first bound follows from a simple iterative algorithm, the second from a modification of [25], and the third from the algorithm of Saks Zhou [24] (as recently improved by Hoza [13]).

1.1 Main Result

We achieve near-optimal space complexity for derandomizing branching programs up to $w = \exp(\log^{1/2} n)$, and improve on known results for all $w \in [\log^{\omega(1)} n, \exp(\log^{.99} n)]$.

THEOREM 1.3. *There is a deterministic algorithm such that given $n, w \in \mathbb{N}$ and a read-once branching program B of length n and width w , the algorithm approximates the acceptance probability of B up to error $1/n^2$. The algorithm runs in space*

$$\tilde{O}\left(\log n + \sqrt{\log n} \cdot \log w\right).$$

This result can be viewed as an attack on BPL vs L from a different direction. While Saks and Zhou (with a further improvement from Hoza) [13, 24] decrease the space required to derandomize a width n , length n branching program, we increase the maximum width w such that we can near-optimally derandomize a width w , length n branching program. This can be thought of as a non-black-box analogue of trying to “build up” better PRGs by starting with the constant width regime [6, 12, 16].

We state our main result in terms of computing an approximate product of stochastic matrices. For this theorem, we without loss of generality assume $w \leq n$.

THEOREM 1.4. *There is an algorithm \mathcal{A}_1 that, given $n, w \in \mathbb{N}$ with $w \leq n$ and arbitrary stochastic matrices $M_1, \dots, M_n \in \mathbb{R}^{w \times w}$, returns a matrix \tilde{M} satisfying*

$$\|M_1 \cdots M_n - \tilde{M}\| \leq 1/n^2.$$

Furthermore, the space complexity of the algorithm is

$$O\left((\log n + \sqrt{\log n} \cdot \log w) \log \log n\right).$$

We note that where we may assume additional structure on M_1, \dots, M_n , such as them being the random walk matrices of undirected or Eulerian graphs, we have existing near-optimal algorithms via a line of work on the Laplacian paradigm [1, 17, 18].

Concurrent Work. Independently and concurrently, Cohen, Doron, Sberlo, and Ta-Shma [11] obtained a matching result to Theorem 1.4, using a different approach (in particular, their algorithm retains the random shift technique).

1.2 Overview of Prior Work

As our work builds on that of Saks and Zhou and Cohen, Doron, and Sberlo, we give a high-level presentation of their approaches before describing our improvements.

The Work of Saks and Zhou [24]. We first present the approach of [24]. For simplicity we assume we are approximating the n th power of a single stochastic matrix $M \in \mathbb{R}^{n \times n}$.

The first step by Saks and Zhou was to divide the problem of computing the n th power of a stochastic matrix M into $\sqrt{\log n}$ iterations of computing the $\ell := 2^{\sqrt{\log n}}$ th power of matrices $M, M^\ell, M^{\ell^2}, \dots$. In each iteration, they used the Nisan generator, which we denote NIS. NIS has an “offline” seed h , which in this case can be of length $O(\log^{3/2} n)$, and an “online” seed of length $O(\log n)$, such that for every fixed offline seed h the generator NIS_h can be evaluated in space $O(\log n)$. Moreover, with high probability over h the generator NIS_h produces a good approximation (in expectation over the *online* seed) of the ℓ th power of a fixed stochastic matrix. They then reused the offline seed throughout the $\sqrt{\log n}$ levels of recursion, at each step using NIS_h to approximate the ℓ th power of the previous level. Thus, the algorithm uses $O(\log n)$ bits of space for the online seed at each of the $\sqrt{\log n}$ levels, and a single offline seed of length $O(\log^{3/2} n)$.

However, there is an issue with this construction as stated. The approximation $\tilde{M} \approx M^\ell$ output by NIS_h can (of course) depend on h . But then the claim that NIS_h is good at approximating powers of M with high probability does not necessarily hold, as we reuse the same offline seed h . To avoid this, the key insights from Saks and Zhou was that between each level of recursion, they could randomly shift and round the matrix \tilde{M} . By doing so, they ensure that with high probability the rounded matrix equals the relevant “true” power, and so NIS_h is good for it with high probability. This results in an additional $O(\log n)$ bits of randomness per level, which is tolerable in their regime.

The Work of Cohen, Doron, and Sberlo [10]. We now discuss the subsequent work of [10], which achieves improved space where $w \ll n$ and all matrices are the same. For simplicity, assume for now that $w = 2^{\sqrt{\log n}}$. Their work contributed the crucial insight that one could use the Nisan generator to approximate $2^{\sqrt{\log n}}$ th powers to accuracy $2^{-c\sqrt{\log n}} \gg 1/n$. This reduced the offline seed to $O(\log n)$, and the online seed to $O(\sqrt{\log n})$.

While this larger error would degrade too much if we let it accumulate over the $\sqrt{\log n}$ levels, they instead used Richardson Iteration to reduce the error back to $1/n$ between each level. Richardson Iteration is a method for improving the accuracy of an approximate power with near-optimal space complexity that has seen several recent applications in the space bounded regime [1, 9, 10, 21].

Unfortunately, their approach seems to require that we are approximately powering a single $w \times w$ matrix, and thus cannot be used to derandomize read-once branching programs of width $w \ll n$ and length n . The reason for this relates to the random shifts used in the construction. Cohen et al. prove that the random shifts can be sampled with $O(\sqrt{\log n} + \log w)$ bits per level. Intuitively, this is because for the i th level, there are w^2 distinct values in the true power M^ℓ , and we must ensure that with high probability our rounding threshold does not lie near one of these values. But if instead there are n distinct base matrices, there could be nw^2 “bad”

values per level and so we must invest $O(\log n)$ bits per level in the shift, leading to an eventual space consumption of $O(\log^{3/2} n)$.²

The other approach to dealing with distinct transition matrices, the one taken in the original paper of Saks and Zhou and other works [1, 9, 21, 24], is that given the $w \times w$ transition matrices M_1, \dots, M_n , first embed them as the off-diagonal elements in a nearly $nw \times nw$ matrix:

$$M = \begin{bmatrix} 0 & M_1 & 0 & \dots & 0 \\ 0 & 0 & \ddots & & \vdots \\ 0 & \dots & 0 & M_{n-1} & 0 \\ 0 & & & 0 & M_n \\ 0 & \dots & & & 0 \end{bmatrix} \quad (2)$$

Then the $(1, n+1)$ block of M^n equals $M_1 \cdots M_n$, so an approximate n th power of M can be used to read off an approximation to the product. However, we are unable to take advantage of this method for computing the product of n distinct $w \times w$ matrices when $w = n^{o(1)}$. This is because this block matrix will necessarily be of size at least $n \times n$, and thus result in space consumption $O(\log^{3/2} n)$.

1.3 Our Approach

We obtain Theorem 1.4 by showing that the random shift step can be *eliminated* from the algorithm of [10]. The reason for random rounding in [10, 24] is the approximate product $\tilde{M} \approx M^\ell$, even at iteration 1, can depend on the offline seed h to the Nisan PRG, and so we cannot say that NIS_h is good for powering this new matrix with high probability. The solution of prior works is to randomly shift \tilde{M} before rounding, such that it is *exactly* equal to (the shifted and rounded version of) M^ℓ , and then apply NIS_h . Since M^ℓ is defined without reference to h , we can say that NIS_h is good for it with high probability.

We take a different approach, by developing a more sophisticated analysis of the error incurred when using NIS_h to approximate powers. Given a stochastic matrix $M \in \mathbb{R}^{w \times w}$, we feed M through a canonicalizer C_t that outputs a branching program B such that the corresponding transition matrix for each layer of B is approximately M . We can then estimate walk probabilities using a suitable pseudorandom generator (in this case, NIS_h). More precisely, from state $v \in [w]$ the canonicalizer assigns the edge with label $\sigma \in [2^t]$ to the state k such that

$$\sum_{j=1}^{k-1} M_{v,j} \leq 2^{-t} \cdot \sigma < \sum_{j=1}^k M_{v,j}.$$

We observe that given M, M' that are close (say, within γ in ℓ_∞ distance), their canonicalizations $B := C_t[M], B' := C_t[M']$ are close *as labeled branching programs*. In particular, from every state v , most of the edges from state v have exactly the same destination in B and B' . Moreover, the only edges that differ must be those assigned in places where the partial sums of rows of M and M' differ. These locations can be (roughly) determined knowing only the partial sums of M and a bound on the ℓ_∞ distance of M' to M . More concretely, there is a set of roughly $2^t \gamma$ “boundary” edges from each state such that for *every* M' such that $\|M - M'\| \leq \gamma/w$,

²They also gave an algorithm based on the Cayley Hamilton theorem that likewise does not seem to generalize to distinct matrices.

every difference between B and $B' := C_t[M']$ will occur on these edges. Note that this set depends only on M and γ , and not on the specific M' . Thus, letting E be the program that accepts if we traverse a boundary edge, a generator NIS_h that fools B and E must also fool $B' = C_t[M']$, even for M' that depend on the offline seed h . We give the formal statement of this result here, and prove it in Section 3:

THEOREM 1.5. *Fix $t \in \mathbb{N}$ and $\gamma \geq 2^{-t}$ and let $M_1, \dots, M_\ell \in \mathbb{R}^{w \times w}$ be sub-stochastic matrices. Suppose $\text{GEN} : \{0, 1\}^d \rightarrow [2^t]^\ell$ is ε -good for $C_t[M_1, \dots, M_\ell]$ and E , where E is defined in Lemma 3.5 only in terms of M_1, \dots, M_ℓ and t and γ . Then for every sub-stochastic $\tilde{M}_1, \dots, \tilde{M}_\ell \in \mathbb{R}^{w \times w}$ where $\|M_i - \tilde{M}_i\| \leq \gamma$ for every i , GEN is $\rho := 6w\ell\gamma + 2\varepsilon$ -good for $C_t[\tilde{M}_1, \dots, \tilde{M}_\ell]$.*

Since an ℓ_∞ error of magnitude γ can shift the value of all w partial sums by γ (and hence a $w \cdot \gamma$ fraction of edges could be allocated differently from each vertex), we obtain the promise that if M, M' are γ close, we fool $C_t[M']$ up to error roughly $\gamma \cdot w$. While this is too much for the original algorithm of Saks and Zhou (except in the small-width regime, see Theorem 1.8), we can use the idea of Cohen, Doron, and Sberlo [10] to reduce the error back to $O(\gamma)$ using Richardson iteration. In fact, we prove that their exact algorithm with the random shift step deleted computes a good approximation. This results in a particularly simple analysis of the final algorithm, which we view as an advantage of our approach.

Remark 1.6. We remark that our canonicalizer can be viewed as truncating the matrix to t bits of precision, then locally monotone-izing the edges from each vertex (in that the state reached from vertex v is a non-decreasing function of the edge label σ , for every v). This local monotone-ization procedure has found several applications in the black box setting, both in analyzing PRGs and bounding the advantage of programs on the coin problem [4, 5, 7, 16]. While previous results used local monotone-ization in the analysis, we take advantage of the non-black-box setting to actually *construct* a locally monotone branching program and apply the PRG on it. In Theorem 1.8, we give a space $O(\log^{3/2} n)$ algorithm for derandomizing programs of width up to $\exp(\log^{1/2} n)$ for which this is the only non-black-box operation in the algorithm.

Decreasing the Failure Probability. We must make one final modification to the algorithm of [10] to handle distinct matrices. They use the Nisan PRG [19] to approximate powers. Our result does not seem to work with the Nisan PRG – we instead use the Nisan PRG combined with the “sampler trick” of Armoni [2]. This is because we require the failure probability δ of the generator to be polynomially small in n (over the offline randomness), not merely w as in [10], because there are nw distinct subprograms to approximate. For Nisan, this would force an *online* seed of length $O(\log n)$, which would require too much space.

To obtain an online seed of length $O(\sqrt{\log n} + \log w)$, we compose Nisan with an averaging sampler with doubly logarithmic dependence on δ in the online component. We note that evaluating the averaging sampler on an online seed (of length $O(\sqrt{\log n} + \log w)$) now takes workspace $O(\log n + \sqrt{\log n} \cdot \log w)$. However, since this workspace is purely used to produce the sampler output (and

hence we do not read the input while using it) we can reuse a fixed $O(\log n + \sqrt{\log n} \cdot \log w)$ bits of workspace for these evaluations across all levels of recursion.

Interestingly, our argument relies on the sampler (and initial PRG) being an *averaging* sampler, because we use a sandwiching argument. The original argument of Saks and Zhou (and Cohen et al.) works more generally for a *weighted* PRG, and this was crucial for Hoza’s improved derandomization of BPL [13]. We view whether our approach works for weighted PRGs to be an interesting open question.

We give a sketch of our algorithm in the case that $M_1 = M_2 = \dots = M_n$. For the formal description, see Algorithm 2.

Algorithm 1: Algorithm sketch

```

1 Draw a random seed  $h$  and let  $\text{NIS}_h$  be the generator with
  offline seed  $h$ .
2 Set  $\tilde{M}^0 = M_1$ .
3 for  $i = 1, \dots, \sqrt{\log n}$  do
4   Canonicalize  $\tilde{M}^{i-1}$  to a degree  $2^{c\sqrt{\log n}}$  branching
     program denoted  $C_{t_1}[\tilde{M}^{i-1}]$ .
5   Use  $\text{NIS}_h$  to approximate walks on  $C_{t_1}[\tilde{M}^{i-1}]^{2^{\sqrt{\log n}}}$  up
     to error  $2^{-c\sqrt{\log n}}$ .
6   Use  $O(\log \log n)$  applications of the space efficient
     Richardson Iteration algorithm to reduce the error of
     the previous matrix to  $1/n^c$  and call the matrix
     resulting from this step  $\tilde{M}^i$ . This reduces the error
     sufficiently such that we can apply Nisan’s generator
     on  $\tilde{M}^i$  despite the fact that this resulting matrix
     depends on  $h$ .
7 end
8 return  $\tilde{M}^{\sqrt{\log n}}$ .
```

1.4 Other Results

Our approach allows us to eliminate the random shifts in the Saks-Zhou algorithm in a range of regimes, which we explicate here. First, by applying Richardson iteration to Theorem 1.4, we can boost our main result to arbitrarily low error:

Corollary 1.7. *Given $w, n \in \mathbb{N}$ with $w \leq n$ and $\varepsilon > 0$ and arbitrary stochastic matrices $M_1, \dots, M_n \in \mathbb{R}^{w \times w}$, there is an algorithm \mathcal{A}'_1 that returns a matrix \tilde{M} satisfying*

$$\|M_1 \cdots M_n - \tilde{M}\| \leq \varepsilon.$$

Furthermore, the space complexity of the algorithm is

$$O\left((\log n + \sqrt{\log n} \cdot \log w) \log \log(n) \log \log(1/\varepsilon) + \log^2 \log(1/\varepsilon)\right).$$

Furthermore, we show that for matrices that are not too wide, the naive approach of iteratively applying the Nisan generator (with $1/\text{poly}(n)$ error) and canonicalizing works, even without Richardson Iteration. This does not improve on Savitch, but we include it as the resulting algorithm is particularly simple. In particular, the only non-black-box step is the canonicalization operation.

THEOREM 1.8. *There is an algorithm \mathcal{A}_2 that, given $n \in \mathbb{N}$ and arbitrary stochastic matrices $M_1, \dots, M_n \in \mathbb{R}^{w \times w}$ where $w \leq 2^{\sqrt{\log n}}$, returns a matrix \tilde{M} satisfying*

$$\|M_1 \cdots M_n - \tilde{M}\| \leq 1/n^2.$$

Furthermore, the space complexity of the algorithm is $O(\log^{3/2} n)$.

Finally, we show that we can achieve space complexity matching Saks-Zhou (with no loglog factors) without random shifts, by interleaving approximate powers with a constant number of Richardson iterations.

THEOREM 1.9. *There is an algorithm \mathcal{A}_3 that, given $n, w \in \mathbb{N}$ and a stochastic matrix $M \in \mathbb{R}^{w \times w}$, returns a matrix \tilde{M} satisfying*

$$\|M^n - \tilde{M}\| \leq 1/w^2.$$

Furthermore, the space complexity of the algorithm is

$$O\left(\log^{3/2} n + \sqrt{\log n} \cdot \log w\right).$$

1.5 Organization

In Section 2 we define terms and recall the Nisan generator (with a modification to decrease its failure probability) and space-efficient Richardson iteration. In Section 3 we formally define the canonicalizer and prove Theorem 1.5. In Section 4 we apply Theorem 1.5 to prove our main result. In Section 5 we show that for small widths, an algorithm without shifts or Richardson iteration can achieve space $O(\log^{3/2} n)$. In Section 6 we show that we can achieve space matching that of Saks and Zhou for large widths, by replacing random shifts by a constant number of Richardson iterations per level.

2 PRELIMINARIES

We first define several terms that we will use in the proofs.

- Let $M \in \mathbb{R}_{\geq 0}^{w \times w}$ be a **stochastic** (resp. **sub-stochastic**) matrix if every row sum is equal (resp. at most) 1. In the language of Markov chains, $M_{i,j}$ is the transition probability from state i to state j .
- Let $\|\cdot\|$ denote the ℓ_∞ matrix norm. We remark that none of our results are sensitive to the precise choice of norm.
- Let U_S denote the uniform distribution over elements of the set S . For $n \in \mathbb{N}$ let U_n be the uniform distribution over $\{0, 1\}^n$.

We recall the property that for sub-stochastic matrices, the product of approximations is an approximation to the product:

Claim 2.1. *Given sub-stochastic $A_1, \dots, A_\ell \in \mathbb{R}^{w \times w}$ and $B_1, \dots, B_\ell \in \mathbb{R}^{w \times w}$, if $\|A_i - B_i\| \leq \delta$ for every i , then*

$$\|A_1 \cdots A_\ell - B_1 \cdots B_\ell\| \leq \ell \cdot \delta.$$

We now formally define branching programs. We define them without reference to a distinguished start or accept state, as we will always require approximations of the walk probability from every state in layer 1 to every state in layer n .

Definition 2.2. A (read-once) **branching program** B of width w and length n with alphabet Σ consists of n functions $B_i : [w] \times$

$[\Sigma] \rightarrow [w]$. We define the composition of branching programs $B := B_1 \cdots B_n$ in the natural way. For $x \in \Sigma^n$ we define

$$\begin{aligned} B[i, x] &= (B_1 \cdots B_n)[i, x] \\ &:= B_n[B_{n-1}[\dots B_2[B_1[i, x_1], x_2] \dots], x_{n-1}, x_n]. \end{aligned}$$

We view the expectation of a branching program as a stochastic matrix. Let $\overline{\mathbb{E}}[B]$ be the matrix where $(\overline{\mathbb{E}}[B])_{i,j} = \Pr[B[i, U_\Sigma^n] = j]$. For a function $\text{GEN} : \{0, 1\}^s \rightarrow \Sigma^n$, let $\overline{\text{GEN}}[B] = \overline{\mathbb{E}}[B \circ \text{GEN}]$.

We say a function $\text{GEN} : \{0, 1\}^s \rightarrow \Sigma^n$ is ϵ -**good** for a branching program B if for every subprogram $B_{i..j}$ we have (truncating the output of GEN to its $j - i$ bit prefix):

$$\left\| \overline{\text{GEN}}[B_{i..j}] - \overline{\mathbb{E}}[B_{i..j}] \right\| \leq \epsilon.$$

Averaging Samplers and Error Reduction Primitives. We recall the family of pseudorandom generators to be used in our construction. We use the Nisan PRG composed with an averaging sampler. We use a formulation of Chattopadhyay and Liao [8].

THEOREM 2.3 ([8, 19, 22]). *Given $n, w, |\Sigma| \in \mathbb{N}$ and $\epsilon, \delta > 0$, there exists a generator $\text{NIS} : \{0, 1\}^m \times \{0, 1\}^d \rightarrow \Sigma^n$ such that for every length n , width w branching program B we have:*

$$\Pr_{h \leftarrow U_m} \left[\left\| \overline{\text{NIS}(h, \cdot)}[B] - \overline{\mathbb{E}}[B] \right\| \leq \epsilon \right] \geq 1 - \delta$$

and we have $m = O(\log(n) \log(nw/\epsilon) + \log(|\Sigma|) + \log(1/\delta))$ and $d = O(\log(nw|\Sigma|/\epsilon) + \log \log(1/\delta))$. Equivalently, for every branching program B , with probability $1 - \delta$ over $h \text{NIS}_h := \text{NIS}(h, \cdot)$ is ϵ -good for B . Furthermore, $\text{NIS}_h(x)$ can be evaluated in space $O(m)$ given two-way read-only access to the offline seed h .

For completeness, we provide a proof of this in Appendix A.

We can transform this generator into a space-efficient algorithm for approximating walks in branching programs, and this is the formulation we will use in the proof:

Lemma 2.4 ([8, 10, 19]). *Given $n, w, |\Sigma| \in \mathbb{N}$ and $\epsilon, \delta > 0$, there exists an algorithm $\text{NIS}_{h,\epsilon,\delta}$ that gets as (read only) input a branching program $B := B_1 \cdots B_n : \Sigma^n \rightarrow \{0, 1\}^{w \times w}$ of width w , accuracy and confidence parameters $\epsilon, \delta > 0$, and $h \in \{0, 1\}^m$ where $m = O(\log(n) \log(nw|\Sigma|/\epsilon) + \log(1/\delta))$. Furthermore:*

- The algorithm runs in space $O(\log(nw|\Sigma|/\epsilon) + \log \log(n/\delta))$, plus an additional $O(m + \log(n/\delta))$ space that is used periodically, during which it does not touch the input or output tapes³.
- The algorithm outputs

$$\left\{ \overline{\text{NIS}}_{i \rightarrow j}[B] \right\}_{i,j \in [n]}$$

where $\overline{\text{NIS}}_{i \rightarrow j}[B]$ is a substochastic matrix. For every $i < j$

and $u, v \in [w]$, $\left(\overline{\text{NIS}}_{i \rightarrow j}[B] \right)_{u,v}$ is the fraction of outputs of

$\text{NIS}_{h,\epsilon,\delta}(x)$ that reach state v in layer j from state u in layer i . If $B_i = B_j$ for all i, j we assume without loss of generality the

algorithm returns $\left\{ \overline{\text{NIS}}_{1 \rightarrow k}[B] \right\}_{k \in [n]}$.

³This space is used for producing the output of the sampler.

- For every branching program B , with probability $1 - \delta$ over $h \leftarrow U_m$ we have that for every $i < j$,

$$\left\| \overline{\text{NIS}}_{i \rightarrow j}[B] - \overline{\mathbb{E}}[B_i \cdots B_j] \right\| \leq \epsilon.$$

We say that $\text{NIS}_{h,\epsilon,\delta}$ is ϵ -**good** for B if this holds.

Likewise, we recall the space-efficient Richardson Iteration algorithm used in prior work.

Lemma 2.5 ([1, 9, 10, 21]). *There exists an algorithm \mathbb{R} that, given sub-stochastic $M_1, \dots, M_n \in \mathbb{R}^{w \times w}$ and $\{\tilde{M}_{i,j}\}_{i < j}$ and $k \in \mathbb{N}$ such that for all $i < j$, $\|M_i \cdots M_j - \tilde{M}_{i,j}\| \leq 1/5n$ returns a sub-stochastic matrix $\mathbb{R}(\{M_i\}_i, \{\tilde{M}_{i,j}\}_{i,j}, k)$ where each entry is represented by at most T bits of precision satisfying*

$$\left\| \mathbb{R}(\{M_i\}_i, \{\tilde{M}_{i,j}\}_{i,j}, k) - M_1 \cdots M_n \right\| \leq 2nw^2 \cdot 2^{-k}.$$

Furthermore, \mathbb{R} runs in space $O(\log^2 k + \log(k) \log(nT))$ where T is the maximum bit complexity of $\{M_i\}, \{\tilde{M}_{i,j}\}$. In the case that $M_i = M$ for all i , we drop the i subscript without loss of generality.

3 THE CANONICALIZER

We first formally define the canonicalizer. Informally, the canonicalizer is an algorithm that converts a substochastic matrix to a branching program, so that we can approximate its expectation using a pseudorandom generator. Our construction crucially relies on how the canonicalizer assigns edges, which we detail below:

Definition 3.1. Given $t \in \mathbb{N}$, there exists a t -**canonicalizer** \mathbb{C}_t that takes in a sub-stochastic matrix $M \in \mathbb{R}^{w \times w}$ with each entry represented by at most s bits and returns a branching program $\mathbb{C}_t[M]$ on $w + 1$ states⁴ with alphabet $\Sigma = [2^t]$. Furthermore, for $i, j, \sigma \in [w] \times [w] \times [2^t]$ the canonicalizer assigns $\mathbb{C}_t[M][i, \sigma] = j$ if and only if

$$\sum_{k=1}^{j-1} M_{i,k} \leq \sigma \cdot 2^{-t} < \sum_{k=1}^j M_{i,k}.$$

Furthermore, \mathbb{C}_t can be computed in space $O(\log(wts))$. We define $\mathbb{C}_t[M_1, \dots, M_\ell] = \mathbb{C}_t[M_1] \cdots \mathbb{C}_t[M_\ell]$ for substochastic matrices M_1, \dots, M_ℓ in the natural way, where now the canonicalizer outputs a branching program of width $w + 1$ and length ℓ and the space complexity is $O(\log(wts\ell))$.

We collect some basic properties of the canonicalizer. In particular, as shown in prior work the expectation of the branching program obtained from canonicalizing M is a close approximation to M .

Claim 3.2 ([10, 24]). *For every $t \in \mathbb{N}$ and sub-stochastic matrix $M \in \mathbb{R}^{w \times w}$, we have*

$$\left\| \overline{\mathbb{E}}[\mathbb{C}_t[M]] - M \right\| \leq w2^{-t}.$$

This extends to canonicalizing a sequence of substochastic matrices:

⁴Since in all cases this additional state will be immaterial, we implicitly pad M to be a (stochastic) $(w+1) \times (w+1)$ matrix and treat the transformation as width preserving.

Claim 3.3. For every $t \in \mathbb{N}$ and sub-stochastic matrices $M_1, \dots, M_\ell \in \mathbb{R}^{w \times w}$, we have

$$\left\| \mathbb{E} [C_t [M_1, \dots, M_\ell]] - M_1 \cdots M_\ell \right\| \leq w\ell 2^{-t}.$$

PROOF. Deferred to Appendix A. \square

We can further extend this to show “canonicalizations of similar matrices are similar”.

Lemma 3.4. Given sub-stochastic M_1, \dots, M_ℓ and $\tilde{M}_1, \dots, \tilde{M}_\ell$ in $\mathbb{R}^{w \times w}$, if $\|M_i - \tilde{M}_i\| \leq \gamma$ for every i , then for every $t \in \mathbb{N}$,

$$\left\| \mathbb{E} [C_t [M_1, \dots, M_\ell]] - \mathbb{E} [C_t [\tilde{M}_1, \dots, \tilde{M}_\ell]] \right\| \leq 2\ell w 2^{-t} + \ell\gamma.$$

PROOF. Deferred to Appendix A. \square

Note that this shows the expectations of the canonicalizations are close, but does not show the canonicalizations are close as labeled branching programs. Proving this stronger result is the primary contribution of the section.

Proof of Theorem 1.5. We now state the main lemma, which proves that canonicalizations of close matrices have similar structures as branching programs. In particular, the edges on which they differ are rare, and can be predicted in advance given only M .

Lemma 3.5. Given $t \in \mathbb{N}$ and sub-stochastic $M_1, \dots, M_\ell \in \mathbb{R}^{w \times w}$ and $\gamma \geq 2^{-t}$, there exists an ordered branching program E of width $w + 2$ and length ℓ and alphabet $[2^t]$ such that:

(1) For every $j \in [w]$,

$$\Pr [E [j, U_{[2^t]^\ell}] = v_{\text{acc}}] \leq 3w\ell \cdot \gamma.$$

(2) For every $\tilde{M}_1, \dots, \tilde{M}_\ell$ satisfying $\|M_i - \tilde{M}_i\| \leq \gamma$ for every i , have that for every state $j \in [w]$ and every input $x \in [2^t]^\ell$,

$$C_t [M_1, \dots, M_\ell] [j, x] \neq C_t [\tilde{M}_1, \dots, \tilde{M}_\ell] [j, x] \implies E [j, x] = v_{\text{acc}}.$$

PROOF. For every $i \in [\ell]$ and state $j \in [w]$, let $s_k := \sum_{l=1}^k (M_i)_{j,l}$ for every $k \in [w]$. Define the boundary set $BD(M_i)_j \subset [2^t]$ of edges for state j in layer i as σ for which

$$BD(M_i)_j := \left\{ \sigma : \sigma \cdot 2^{-t} \in \bigcup_{k=1}^w [s_k - \gamma, s_k + \gamma] \right\}.$$

Let E be the program (with $w + 2$ states) that has identical states and transitions to $C_t [M_1, \dots, M_\ell]$ except it transitions to state v_{acc} if and only if the input ever traverses such an edge (and once it reaches v_{acc} always stays at this state). Observe that for every state j and every layer i , a random edge from j lies in the boundary set with low probability. In particular,

$$\Pr_{\sigma \leftarrow U_{[2^t]}} [C_t [M_i] [j, \sigma] \in BD(M_i)_j] \leq 2\gamma \cdot w + \frac{w}{2^t} \leq 3\gamma w$$

and so a union bound over steps proves (1).

Now fix $\tilde{M}_1, \dots, \tilde{M}_\ell$ satisfying $\|M_i - \tilde{M}_i\| \leq \gamma$ for every i . Observe that the boundary sets and error test program have been defined without reference to \tilde{M} .

Claim 3.6. For every $i \in [\ell]$ and $j \in [w]$, $BD(M_i)_j$ contains all edges σ on which $C_t [M_i] [j, \sigma] \neq C_t [\tilde{M}_i] [j, \sigma]$.

PROOF. Fixing i, j , for each $k \in [w]$ let

$$\tilde{s}_k := \sum_{l=1}^k (\tilde{M}_i)_{j,l},$$

and recall s_k is defined analogously. We have that for every k ,

$$\begin{aligned} |s_k - \tilde{s}_k| &= \left| \sum_{l=1}^k (M_i)_{j,l} - \sum_{l=1}^k (\tilde{M}_i)_{j,l} \right| \\ &\leq \sum_{l=1}^k |(M_i)_{j,l} - (\tilde{M}_i)_{j,l}| \\ &\leq \left\| \tilde{M}_i - M_i \right\| \leq \gamma. \end{aligned}$$

Thus for every k ,

$$[s_k, \tilde{s}_k] \subset [s_k - \gamma, s_k + \gamma].$$

By definition of the canonicalizer, for every σ such that

$$C_t [M_i] [j, \sigma] \neq C_t [\tilde{M}_i] [j, \sigma],$$

we must have $\sigma \cdot 2^{-t} \in [s_k, \tilde{s}_k] \subset [s_k - \gamma, s_k + \gamma]$ for some k , and so $\sigma \in BD(M_i)_j$ as claimed. \square

Then we observe that for every j, x such that

$$C_t [M_1, \dots, M_\ell] [j, x] \neq C_t [\tilde{M}_1, \dots, \tilde{M}_\ell] [j, x],$$

there is i such that

$$v := C_t [M_1, \dots, M_i] [j, x_{1..i}] = C_t [\tilde{M}_1, \dots, \tilde{M}_i] [j, x_{1..i}]$$

and $C_t [\tilde{M}_{i+1}] [v, x_{i+1}] \neq C_t [M_{i+1}] [v, x_{i+1}]$, since otherwise x would make identical transitions in both programs. But then the edge labeled x_{i+1} from v must have received different labels in $C_t [M_{i+1}]$ and $C_t [\tilde{M}_{i+1}]$ and thus lie in $BD(M_{i+1})_j$ by Claim 3.6, and so by definition of E we must have $E [j, x] = v_{\text{acc}}$. Note that it could have been the case that x transited a boundary edge before layer $i + 1$, but this likewise causes $E [j, x]$ to reach state v_{acc} on input x . \square

We can then prove Theorem 1.5 using Lemma 3.5.

THEOREM 1.5. Fix $t \in \mathbb{N}$ and $\gamma \geq 2^{-t}$ and let $M_1, \dots, M_\ell \in \mathbb{R}^{w \times w}$ be sub-stochastic matrices. Suppose $\text{GEN} : \{0, 1\}^d \rightarrow [2^t]^\ell$ is ε -good for $C_t [M_1, \dots, M_\ell]$ and E , where E is defined in Lemma 3.5 only in terms of M_1, \dots, M_ℓ and t and γ . Then for every sub-stochastic $\tilde{M}_1, \dots, \tilde{M}_\ell \in \mathbb{R}^{w \times w}$ where $\|M_i - \tilde{M}_i\| \leq \gamma$ for every i , GEN is $\rho := 6w\ell\gamma + 2\varepsilon$ -good for $C_t [\tilde{M}_1, \dots, \tilde{M}_\ell]$.

PROOF. By Lemma 3.5 applied with $\gamma = \gamma$ and $t = t$ we have

$$\begin{aligned} & \left\| \overline{\text{GEN}} \left[\text{C}_t [M_1, \dots, M_\ell] \right] - \overline{\text{GEN}} \left[\text{C}_t [\tilde{M}_1, \dots, \tilde{M}_\ell] \right] \right\| \\ & \leq \max_j \left\{ \Pr_{x \leftarrow U_{\{0,1\}^d}} \left[\text{C}_t [M_1, \dots, M_\ell] [j, \text{GEN}(x)] \neq \right. \right. \\ & \quad \left. \left. \text{C}_t [M_1, \dots, M_\ell] [j, \text{GEN}(x)] \right] \right\} \\ & \leq \max_j \left\{ \Pr_{x \leftarrow U_{\{0,1\}^d}} \left[E[j, \text{GEN}(x)] = v_{\text{acc}} \right] \right\} \quad (\text{Property 2}) \\ & \leq 3w\ell \cdot \gamma + \varepsilon \end{aligned}$$

where the final line follows from Property 1 and GEN being ε -good for E . Furthermore, since by assumption GEN is ε -good for $\text{C}_t [M_1, \dots, M_\ell]$:

$$\left\| \overline{\text{GEN}} \left[\text{C}_t [M_1, \dots, M_\ell] \right] - \overline{\mathbb{E}} \left[\text{C}_t [M_1, \dots, M_\ell] \right] \right\| \leq \varepsilon$$

and by Lemma 3.4,

$$\begin{aligned} & \left\| \overline{\mathbb{E}} \left[\text{C}_t [M_1, \dots, M_\ell] \right] - \overline{\mathbb{E}} \left[\text{C}_t [\tilde{M}_1, \dots, \tilde{M}_\ell] \right] \right\| \\ & \leq 2\ell w 2^{-t} + \ell \gamma \\ & \leq 3\ell w \gamma. \end{aligned}$$

Thus, applying the triangle inequality we conclude that

$$\left\| \overline{\text{GEN}} \left[\text{C}_t [\tilde{M}_1, \dots, \tilde{M}_\ell] \right] - \overline{\mathbb{E}} \left[\text{C}_t [\tilde{M}_1, \dots, \tilde{M}_\ell] \right] \right\| \leq 6w\ell\gamma + 2\varepsilon. \quad \square$$

We note that in the regime $w = n$, applying this result directly in the original framework of Saks and Zhou [24] does not allow us to eliminate the random rounding step. This is because our error degrades with a factor of w per application, which could give a final error of $w^{\sqrt{\log n}} = n^{\omega(1)}$. However, we can use the approach of Cohen et al. to repair the loss at each level.

4 PROOF OF THEOREM 1.4

We now apply Theorem 1.5 to prove Theorem 1.4. The analysis of our resulting algorithm is cleaner than prior approaches due to the absence of random rounding. In particular, we directly argue that the approximation at level i is close to the $2^i \cdot \sqrt{\log n}$ th true power, rather than comparing to a shifted and rounded version of such. As we make the same parameter choices⁵ as Cohen, Doron, and Sberlo (and the components of our algorithm are a strict subset of theirs), the space complexity follows essentially from their analysis, though we must be careful to avoid incurring an overhead of $O(\log n)$ bits per level for tracking indices or evaluating the sampler. We analyze the space complexity in Lemma 4.3.

We are now prepared to analyze the algorithm. Given n, w as specified in the upcoming theorem (where we may assume $n \geq w$ without loss of generality, as otherwise apply Theorem 1.9), we set parameters as follows:

$$\begin{aligned} r_1 = r_2 = \sqrt{\log n}, \quad t_2 = 20 \log n, \quad t_1 = 2r_1 + \log w + 4 \\ \varepsilon = 2^{-r_1-4}, \quad \delta = 2n^{-5}, \quad R_1 = 2^{r_1}. \end{aligned}$$

⁵With the exception of δ , which we must take to be order $1/n$ to survive a bound over the test programs.

Let $\text{NIS}_{h,\varepsilon,\delta}$ be the family of online-offline samplers from Lemma 2.4 with $n = R_1, w = w + 2, \Sigma = [2^{t_1}], \varepsilon = \varepsilon$ and $\delta = \delta$. We remark that the offline seed has length

$$\begin{aligned} m &= O\left(r_1 \log(R_1 w 2^{t_1}/\varepsilon) + \log(1/\delta) + \log(1/\varepsilon)\right) \\ &= O\left(\log n + \sqrt{\log n} \cdot \log w\right). \end{aligned}$$

Furthermore, the online seed has length

$$d = O\left(\log(R_1 w 2^{t_1}/\varepsilon) + \log \log(1/\delta)\right) = O\left(\sqrt{\log n} + \log w\right).$$

We give the formal description of the algorithm in Algorithm 2.

Algorithm 2: $\text{SZ}(M_1, \dots, M_n, h)$

- 1 Given $h \in \{0, 1\}^m$, let $\text{NIS} := \text{NIS}_{h,\varepsilon,\delta}$ be the sampler with parameters set above.
 - 2 **return** $\text{IMM}(\sqrt{\log n}, 0, \text{NIS})$.
-

Algorithm 3: $\text{IMM}(i, j, \text{NIS})$

- 1 **if** $i = 0$ **then**
 - 2 | **return** M_j^0 ;
 - 3 **end**
 - 4 **for** $l \in [j, j + R_1]$ **do**
 - 5 | Let $\tilde{M}_l^{i-1} := \text{IMM}(i-1, R_1 \cdot j + l, \text{NIS})$.
 - 6 **end**
 - 7 Let

$$\left\{ \tilde{M}_{k,l} \right\}_{k,l \in [R_1]} := \left\{ \overline{\text{NIS}}_{k \rightarrow l} \left[\text{C}_{t_1} \left[\tilde{M}_j^{i-1}, \dots, \tilde{M}_{j+R_1}^{i-1} \right] \right] \right\}_{k,l \in [R_1]}$$
 - 8 Set

$$\tilde{M}_j^i := \mathbb{R} \left(\left\{ \tilde{M}_{j+l}^{i-1} \right\}_{l \in [R_1]}, \left\{ \tilde{M}_{k,l} \right\}_{k,l \in [R_1]}, t_2 \right)$$
 - 9 **return** \tilde{M}_j^i .
-

We now show Algorithm 2 works with high probability over the outer seed h .

THEOREM 4.1. *Given $n, w \in \mathbb{N}$ with $w \leq n$ and arbitrary sub-stochastic matrices $M_1, \dots, M_n \in \mathbb{R}^{w \times w}$, Algorithm 2 returns with probability $1 - n^2\delta$ over the outer seed $h \leftarrow \{0, 1\}^m$ a matrix $\tilde{M} = \text{SZ}(M_1, \dots, M_n, h)$ satisfying*

$$\left\| M_1 \cdots M_n - \tilde{M} \right\| \leq 1/n^3.$$

Moreover, $\text{SZ}(M_1, \dots, M_n, h)$ runs in space

$$O\left(\left(\log n + \sqrt{\log n} \cdot \log w\right) \log \log(n)\right)$$

We first define the true powers that we wish to approximate, and their corresponding canonicalizations.

Definition 4.2. Given $n, w, t_1 \in \mathbb{N}$ and sub-stochastic $M_1^0, \dots, M_n^0 \in \mathbb{R}^{w \times w}$, define

$$M_j^i = M_{j-R_1}^{i-1} \cdot M_{j-R_1+1}^{i-1} \cdots M_{j-R_1+1}^{i-1}$$

and let $\mathcal{P} = \left\{ \mathcal{C}_{t_1} \left[M_j^i, \dots, M_{j+R_1}^i \right] \right\} \cup \mathcal{E}$ where \mathcal{E} is the family of error testers of Lemma 3.5 applied with $t = t_1$ and $\gamma = \gamma$ (for γ to be globally chosen later) to every subproduct $M_j^i, \dots, M_{j+R_1}^i$.

Note that we have $M_1^{r_2} = M_1^0 \cdots M_n^0$ by definition. We first prove the correctness, then analyze the space consumption.

PROOF OF CORRECTNESS OF THEOREM 4.1. We condition on the event that $\text{NIS} := \text{NIS}_{h,\varepsilon,\delta}$ is ε -good for the set of programs \mathcal{P} as defined in Definition 4.2. This occurs with probability at least $1 - n^2\delta$ by Lemma 2.4 and the fact that there are at most $nw \leq n^2$ such programs. Subsequent to this assumption (which requires a union bound over nw bad events, rather than w), the proof does not change if we assume all base matrices are equal, so we do so for clarity. We maintain the following invariant at level i of the algorithm:

$$\left\| \tilde{M}^i - M^i \right\| \leq \frac{2^{2i \cdot r_1}}{n^{10}}$$

Ensuring this invariant holds certainly suffices to complete the proof. Assuming the invariant holds for level i , we now verify that the conditions of Theorem 1.5 are satisfied for $t = t_1$ and $\ell = R_1$ and $\varepsilon = \varepsilon$ and $\gamma := 2^{-t_1}$ and $M = M^i$ and $\tilde{M} = \tilde{M}^i$. We have that NIS is ε -good for M^i and the associated error tester by assumption. Furthermore by the invariant we have

$$\left\| \tilde{M}^i - M^i \right\| \leq \frac{1}{n^2} \leq \gamma.$$

Therefore by Theorem 1.5 applied to the generator NIS , we obtain for every $j \in [R_1]$,

$$\begin{aligned} \left\| \mathbb{E} \left[\mathcal{C}_{t_1} \left[\tilde{M}^i, \dots, \tilde{M}^i \right] \right] - \overline{\text{NIS}}_{1 \rightarrow j} \left[\mathcal{C}_{t_1} \left[\tilde{M}^i, \dots, \tilde{M}^i \right] \right] \right\| &\leq 6wR_1\gamma + 2\varepsilon \\ &\leq \frac{1}{10 \cdot R_1} \end{aligned}$$

And thus by Claim 3.3 for every $j \in [R_1]$,

$$\left\| (\tilde{M}^i)^j - \overline{\text{NIS}}_{1 \rightarrow j} \left[\mathcal{C}_{t_1} \left[\tilde{M}^i, \dots, \tilde{M}^i \right] \right] \right\| \leq \frac{1}{10 \cdot R_1} + wR_1 2^{-t_1} \leq \frac{1}{5 \cdot R_1}.$$

Therefore, recalling

$$\tilde{M}^{i+1} := \mathbb{R} \left(\tilde{M}^i, \left\{ \overline{\text{NIS}}_{1 \rightarrow j} \left[\mathcal{C}_{t_1} \left[\tilde{M}^i, \dots, \tilde{M}^i \right] \right] \right\}_j, t_2 \right)$$

by Lemma 2.5 we have

$$\left\| \tilde{M}^{i+1} - (\tilde{M}^i)^{R_1} \right\| \leq 2 \cdot R_1 w^2 2^{-t_2} \leq \frac{1}{n^{10}}. \quad (3)$$

Thus,

$$\begin{aligned} \left\| \tilde{M}^{i+1} - M^{i+1} \right\| &\leq \left\| \tilde{M}^{i+1} - (\tilde{M}^i)^{R_1} \right\| + \left\| (\tilde{M}^i)^{R_1} - M^{i+1} \right\| \\ &\leq \frac{1}{n^{10}} + \left\| (\tilde{M}^i)^{R_1} - (M^i)^{R_1} \right\| \\ &\leq \frac{2}{n^{10}} + R_1 \cdot \left\| \tilde{M}^i - M^i \right\| \\ &\leq \frac{2^{2(i+1) \cdot r_1}}{n^{10}} \end{aligned} \quad \text{Claim 2.1}$$

which maintains the invariant for the next level.

Note that since that $i \leq \sqrt{\log n}$, this means that our error is bounded with

$$\frac{2^{2i \cdot r_1}}{n^{10}} \leq \frac{2^{2\sqrt{\log n} \cdot \sqrt{\log n}}}{n^{10}} \leq \frac{1}{n^8}. \quad \square$$

Lemma 4.3. $\text{SZ}(M_1, \dots, M_n, h)$ runs in space

$$O \left((\log n + \sqrt{\log n} \cdot \log(w)) \log \log(n) \right).$$

PROOF. We note that at no point do we explicitly write down the matrix \tilde{M}_j^i , which would require $w^2 \log(nw)$ bits. Instead, whenever $\text{IMM}(i+1, j)$ requests a bit of \tilde{M}_j^i , we recurse on $\text{IMM}(i, j')$ and determine only this bit, then return control to level $i+1$. This process is formalized as the composition of space-bounded algorithms in Lemma A.2.

We look at the individual space complexities of the components of our algorithm:

- (1) First, we note that the seed length for the generator requires space (paid once)

$$O \left(\log n + \sqrt{\log n} \cdot \log w \right).$$

- (2) Each function $\text{IMM}(i, j)$ produces $w^2 t_2 = O(w^2 \log n)$ bits of output, and so by Lemma A.2 we require $O(\log w + \log \log n)$ bits per level to track the index of the bit to be output.
- (3) The online space for NIS requires space (paid once per level)

$$O \left(\sqrt{\log n} + \log w \right).$$

- (4) By Lemma 2.4, we require space $O(\log n + \sqrt{\log n} \cdot \log w)$ to evaluate $\text{NIS}_h(x)$ on online input x , and as we do not touch the input or output during this time, this space can be reused between levels and so only needs to be paid for once.
- (5) Richardson iteration requires space (paid once per level)

$$O \left(\log^2 t_2 + \log t_2 \cdot \log(R_1 w t_1) \right)$$

$$= O \left((\log \log n)^2 + \log \log n \cdot (\log w + \sqrt{\log n}) \right).$$

- (6) The canonicalizer requires space (paid once per level)

$$O(\log(t_1 t_2 w R_1)) = O(\sqrt{\log n} + \log w + \log \log n).$$

- (7) Specifying which matrices on the input tape should be multiplied can be done with $O(\log R_1) = O(\sqrt{\log n})$ bits per level of recursion.

To justify the last statement above, we see that if we assume the algorithm stores the index of which subproblem is currently being solved in each level, it can directly compute which matrices are supposed to be multiplied at any given step. Each index takes space $O(\log R_1)$, and is stored only once for each level in the call tree. Bringing this together, we note that the maximum recursion depth of our algorithm is $r_2 = \sqrt{\log n}$, and the space complexity of each recursive level is $O(\log w \log \log n + \sqrt{\log n} \log \log n)$. This means in total, the space required for the call tree is

$$O \left(\sqrt{\log n} \cdot \log w \cdot \log \log n + \log n \cdot \log \log n \right),$$

and this complexity does not change when we account for the offline seed of Nisan's generator. \square

Finally, we can use Theorem 4.1 to prove Theorem 1.4. We note that Cohen et al. [9] obtain a $1/n$ approximation by taking the median of each entry over the offline randomness (as their algorithm fails with probability $1/w \gg 1/n$ over h), but we obtain failure probability $1/n$ over h , so we take the average for simplicity.

PROOF OF THEOREM 1.4. By our choice of δ , with probability at least $1 - 1/2n^3$ over the outer seed we obtain a final matrix \tilde{M} satisfying $\|\tilde{M} - M_1 \cdots M_n\| \leq n^{-3}$. For a bad h , we receive a matrix with distance at most $w \leq n$ in ℓ_∞ distance, and so the theorem follows from letting \mathcal{A}_1 be the algorithm that returns the average of Algorithm 2 over h . \square

By applying a final layer of Richardson iteration, we can obtain an arbitrary low-accuracy estimate at mild additional cost in space. We use a more precise statement of Richardson iteration, but we defer its statement to Section 6, where we use it as part of the inner loop of an algorithm.

PROOF OF THEOREM 1.7. Let Algorithm \mathcal{A}'_1 be the algorithm that applies Lemma 6.1 to M_1, \dots, M_n and $\{\tilde{M}_{i,j}\}_{i,j \in [n]}$ with error ε , where for i, j , $\tilde{M}_{i,j}$ is the output of Theorem 1.4 applied to M_i, \dots, M_j . The correctness is direct from the correctness of Theorem 1.4, and the space complexity follows from Lemma 6.1 and the composition of space bounded algorithms (Lemma A.2). \square

5 NAIVE SAKS-ZHOU FOR SMALL WIDTH

In this section we prove that the naive Saks-Zhou algorithm succeeds without random shifts as long as $w = O(2^{\sqrt{\log n}})$. We remark that this gives a space $O(\log^{3/2} n)$ algorithm for derandomizing width $\exp(\log^{1/2} n)$ branching programs whose only non-black-box step is repeatedly locally monotonizing subprograms (see Remark 1.6). For the remainder of the section, let

$$t = 20 \log n, \quad \varepsilon = n^{-10}, \quad \delta = n^{-6}, \quad r_1 = r_2 = \sqrt{\log n}, \quad R_1 = 2^{t_1}.$$

Let $\text{NIS}_{h,\varepsilon,\delta}$ be the family of online-offline samplers from Lemma 2.4 with $n = R_1$, $w = w + 2$, $\Sigma = [2^t]$, $\varepsilon = \varepsilon$ and $\delta = \delta$. We remark that the offline seed has length

$$m = O(r_1 \log(r_1 w 2^t / \varepsilon) + \log(1/\delta) + \log(1/\varepsilon)) = O(\log^{3/2} n).$$

Furthermore, the online seed has length

$$d = O(\log(r_1 w 2^t / \varepsilon) + \log \log(1/\delta)) = O(\log n).$$

We now formally state this algorithm.

Algorithm 4: $\text{SZN}(M_1, \dots, M_n, h)$

- 1 Given $h \in \{0, 1\}^m$, let $\text{NIS} := \text{NIS}_{h,\varepsilon,\delta}$ be the sampler with parameters set above.
 - 2 **return** $\text{NaiveIMM}(\sqrt{\log n}, 0, \text{NIS})$.
-

We then state the theorem showing it is correct with high probability over the outer seed:

Algorithm 5: Algorithm $\text{NaiveIMM}(i, j, \text{NIS})$

- 1 **if** $i = 0$ **then**
 - 2 | **return** M_j^0
 - 3 **end**
 - 4 **for** $l \in [j, j + R_1]$ **do**
 - 5 | let $\tilde{M}_l^{i-1} := \text{NaiveIMM}(i-1, R_1 \cdot j + l, \text{NIS})$.
 - 6 **end**
 - 7 **return** $\overline{\text{NIS}}_{1 \rightarrow R_1} \left[\text{C}_{t_1} \left[\tilde{M}_j^{i-1}, \dots, \tilde{M}_{j+R_1}^{i-1} \right] \right]$.
-

THEOREM 5.1. Given $n \in \mathbb{N}$ and $w \leq 2^{\sqrt{\log n}}$ and any sub-stochastic $M_1, \dots, M_n \in \mathbb{R}^{w \times w}$, Algorithm 4 returns with probability $1 - n^2 \delta$ over the outer seed $h \leftarrow \{0, 1\}^m$ a matrix $\tilde{M} = \text{SZN}(M_1, \dots, M_n, r_1, r_2, h)$ satisfying

$$\|M_1 \cdots M_n - \tilde{M}\| \leq 1/n^3.$$

Moreover, $\text{SZN}(M_1, \dots, M_n, h)$ runs in space $O(\log^{3/2} n)$.

We construct the true powers identically to Definition 4.2 (with $t_1 = t$ and $\gamma_1, \dots, \gamma_{r_2}$ to be chosen later, where the choice depends on the recursion level and nothing else). Furthermore, the space complexity directly follows from the analysis in the prior section, where we no longer pay for Richardson iterations. The remainder of the proof consists of setting parameters and walking through the same analysis.

PROOF OF THEOREM 5.1. We condition on the event that $\text{NIS} := \text{NIS}_h$ is ε -good for \mathcal{P} as defined in Definition 4.2, which occurs with probability at least $1 - n^2 \delta$ as claimed. As in the prior case, subsequent to this assumption the proof does not change if we assume all base matrices are equal, so we do so for clarity. We maintain the following invariant at the start of the i th iteration of the loop:

$$\|\tilde{M}^i - M^i\| \leq \frac{(8wR_1)^i R_1^i}{n^{10}}$$

Ensuring this invariant holds certainly suffices to complete the proof. We now verify that the conditions of Theorem 1.5 are satisfied for $t = t$ and $\varepsilon = \varepsilon$ and $\gamma_i := (8wR_1)^i R_1^i / n^{10}$ and $M = M^i$ and $\tilde{M} = \tilde{M}^i$. We have that NIS is ε -good for the relevant programs by assumption. Furthermore by the invariant we have

$$\|\tilde{M}^i - M^i\| \leq \gamma_i.$$

Therefore by Theorem 1.5, recalling that

$$\tilde{M}^{i+1} := \overline{\text{NIS}}_{1 \rightarrow R_1} \left[\text{C}_t \left[\tilde{M}^i, \dots, \tilde{M}^i \right] \right]$$

we obtain

$$\|\mathbb{E} \left[\text{C}_t \left[\tilde{M}^i, \dots, \tilde{M}^i \right] \right] - \tilde{M}^{i+1}\| \leq 6wR_1 \gamma_i + 2\varepsilon \leq (7wR_1) \gamma_i.$$

and thus by Claim 3.3,

$$\|(\tilde{M}^i)^{R_1} - \tilde{M}^{i+1}\| \leq (7wR_1) \gamma_i + wR_1 2^{-t} \leq (8wR_1) \gamma_i. \quad (4)$$

Thus,

$$\begin{aligned}
\left\| \tilde{M}^{i+1} - M^{i+1} \right\| &\leq \left\| \tilde{M}^{i+1} - (\tilde{M}^i)^{R_1} \right\| + \left\| (\tilde{M}^i)^{R_1} - M^{i+1} \right\| \\
&\leq (8wR_1)\gamma_i + \left\| (\tilde{M}^i)^{R_1} - (M^i)^{R_1} \right\| \quad (4) \\
&\leq (8wR_1)\gamma_i + R_1 \cdot \left\| \tilde{M}^i - M^i \right\| \quad \text{Claim 2.1} \\
&\leq \frac{(8wR_1)^{i+1} R_1^{i+1}}{n^{10}} + \frac{(8wR_1)^i R_1^{i+1}}{n^{10}} \\
&\leq \frac{(8wR_1)^{i+1} R_1^{i+1}}{n^{10}}.
\end{aligned}$$

which maintains the invariant for the next level. Note that since that $i \leq \sqrt{\log n}$, this means that our error is bounded with

$$\frac{(8wR_1)^i R_1^i}{n^{10}} \leq \frac{n^2 \cdot (8w)^{\sqrt{\log n}}}{n^{10}} = \frac{(8w)^{\sqrt{\log n}}}{n^8}. \quad \square$$

Hence, we see that when $w \leq 2\sqrt{\log n}$, the original algorithm of Saks and Zhou works without random shifting. The proof of Theorem 1.8 from Theorem 5.1 is exactly analogous to that of Theorem 1.4 from Theorem 4.1, so we omit it.

6 SAKS-ZHOU WITH CONSTANT RICHARDSON ITERATIONS

To obtain space complexity matching Saks-Zhou without random rounding (with no loglog factors), we state a more precise version of Richardson iteration that utilizes stronger guarantees on the initial error:

Lemma 6.1 ([1, 9, 10, 21]). *There exists an algorithm R that, given $t, k \in \mathbb{N}$ and sub-stochastic $M \in \mathbb{R}^{w \times w}$ and $(\tilde{M}_i)_{i \in [n]}$ and $k \in \mathbb{N}$ such that for all i $\|M^i - \tilde{M}_i\| \leq 1/(nw)^2$ returns a substochastic matrix $R(M, (\tilde{M}_i)_i, t, k)$ where each entry is represented by at most t bits of precision satisfying*

$$\left\| R\left(M, \{\tilde{M}_i\}_i, t, k\right) - M^n \right\| \leq (nw)^{-k}.$$

Furthermore, R runs in space $O(\log^2 k + \log(k) \log(tnT))$ where T is the maximum bit complexity of M and $\{M_i\}$.

Note that the above lemma implies that with only $O(1)$ Richardson iterations, we can reduce the error from $\frac{1}{w^c}$ back down to $\frac{1}{w^{c'}}$ for arbitrary constants $c < c'$, and this consumes $O(\log nw)$ space per level.

As we are now in the regime where we may assume $w \geq n$ without loss of generality, we appeal to identity 2 and assume we are computing the n th power of a single $w \times w$ stochastic matrix. For the remainder of the section let

$$t = 20 \log(w), \varepsilon = w^{-10}, \delta = w^{-6}, r_1 = r_2 = \sqrt{\log n}, R_1 = 2^{r_1}.$$

Let $\text{NIS}_{h,\varepsilon,\delta}$ be the family of online-offline samplers from Lemma 2.4 with $n = R_1$, $w = w + 2$, $\Sigma = [2^t]$, $\varepsilon = \varepsilon$ and $\delta = \delta$. We remark that the offline seed has length

$$\begin{aligned}
m &= O\left(r_1 \log(R_1 w 2^t / \varepsilon) + \log(1/\delta) + \log(1/\varepsilon)\right) \\
&= O\left(\log^{3/2} n + \sqrt{\log n} \cdot \log w\right).
\end{aligned}$$

Furthermore, the online seed has length

$$d = O\left(\log(R_1 w 2^t / \varepsilon) + \log \log(1/\delta)\right) = O(\log nw).$$

We formally describe the inner loop as Algorithm 6.

Algorithm 6: Algorithm SZC(M, h)

- 1 Given $h \in \{0, 1\}^m$, let $\text{NIS}_{h,\varepsilon,\delta}$ be the sampler with parameters set above.
- 2 Set $\tilde{M}^0 := M$.
- 3 **for** $i = 1, \dots, r_2$ **do**
- 4 Let

$$\left\{ \tilde{M}_k \right\}_{k \in [R_1]} := \left\{ \text{NIS}_{1 \rightarrow k} \left[\text{Ct}_1 \left[\tilde{M}_1^{i-1}, \dots, \tilde{M}_{R_1}^{i-1} \right] \right] \right\}_{k \in [R_1]}.$$
- 5 Set

$$\tilde{M}^i := R\left(\tilde{M}^{i-1}, \left\{ \tilde{M}_k \right\}_{k \in [R_1]}, t, 10\right).$$
- 6 **end**
- 7 **return** \tilde{M}^{r_2}

We then show Algorithm 6 is correct with high probability over the outer seed:

THEOREM 6.2. *Given $n, w \in \mathbb{N}$ with $w \geq n$ and any stochastic $M \in \mathbb{R}^{w \times w}$, Algorithm 6 returns with probability $1 - w^{-2}$ over the outer seed $h \leftarrow \{0, 1\}^m$ a matrix $\tilde{M} = \text{SZC}(M, h)$ satisfying*

$$\left\| M^n - \tilde{M} \right\| \leq 1/w^3.$$

Moreover, $\text{SZC}(M, h)$ runs in space

$$O\left(\log^{3/2} n + \sqrt{\log n} \cdot \log w\right).$$

We construct the true powers identically to Definition 4.2 (with $t_1 = t$). The rest of the proof is analogous to Theorem 4.1, except that we set γ to be $1/w^c$ and so require only a constant number of Richardson iterations per level. The space complexity follows from Lemma 6.1, as we now require $O(\log nw)$ space per level for the Richardson iterations (and the online seed).

ACKNOWLEDGEMENTS

We thank Salil Vadhan and David Zuckerman for their insightful comments during the writing of this paper.

A DEFERRED PROOFS

We collect proofs of claims regarding the accuracy of various truncation and rounding procedures. In all cases, our results are insensitive to polynomial losses in the length and width of the relevant sub-programs, and constant factors in terms of the bit complexity (i.e. the parameters t_1, t_2).

We first extend Claim 3.2 to canonicalizations of sequences of matrices.

Claim 3.3. *For every $t \in \mathbb{N}$ and sub-stochastic matrices $M_1, \dots, M_\ell \in \mathbb{R}^{w \times w}$, we have*

$$\left\| \mathbb{E} \left[\text{Ct} \left[M_1, \dots, M_\ell \right] \right] - M_1 \cdots M_\ell \right\| \leq w \ell 2^{-t}.$$

PROOF. We have

$$\begin{aligned} & \left\| \overline{\mathbb{E}} [C_t [M_1, \dots, M_\ell]] - M_1 \cdots M_\ell \right\| \\ &= \left\| \overline{\mathbb{E}} [C_t [M_1]] \cdots \overline{\mathbb{E}} [C_t [M_\ell]] - M_1 \cdots M_\ell \right\| \\ &\leq \ell w 2^{-t} \end{aligned}$$

where the second step uses Claim 2.1 with $A_i = \overline{\mathbb{E}} [C_t [M_i]]$ and $B_i = M_i$ and $\delta = w 2^{-t}$ from Claim 3.2. \square

We observe that this easily implies that the expectations of canonicalizations of similar matrices are similar.

Lemma 3.4. *Given sub-stochastic M_1, \dots, M_ℓ and $\tilde{M}_1, \dots, \tilde{M}_\ell$ in $\mathbb{R}^{w \times w}$, if $\|M_i - \tilde{M}_i\| \leq \gamma$ for every i , then for every $t \in \mathbb{N}$,*

$$\left\| \overline{\mathbb{E}} [C_t [M_1, \dots, M_\ell]] - \overline{\mathbb{E}} [C_t [\tilde{M}_1, \dots, \tilde{M}_\ell]] \right\| \leq 2\ell w 2^{-t} + \ell \gamma.$$

PROOF. We have

$$\begin{aligned} & \left\| \overline{\mathbb{E}} [C_t [M_1, \dots, M_\ell]] - \overline{\mathbb{E}} [C_t [\tilde{M}_1, \dots, \tilde{M}_\ell]] \right\| \\ &\leq \left\| \overline{\mathbb{E}} [C_t [M_1, \dots, M_\ell]] - M_1 \cdots M_\ell \right\| \\ &\quad + \left\| M_1 \cdots M_\ell - \tilde{M}_1 \cdots \tilde{M}_\ell \right\| \\ &\quad + \left\| \tilde{M}_1 \cdots \tilde{M}_\ell - \overline{\mathbb{E}} [C_t [\tilde{M}_1, \dots, \tilde{M}_\ell]] \right\| \\ &\leq 2\ell w 2^{-t} + \ell \gamma \end{aligned}$$

where the final line uses Claim 3.3 and Claim 2.1. \square

We recall the formal statement of the composition of space-bounded algorithms:

Lemma A.1 ([10]). *Let $f_1, f_2 : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be computable in space $s_1, s_2 : \mathbb{N} \rightarrow \mathbb{N}$, where $s_1(n), s_2(n) \geq \log n$. Then, $f_1 \circ f_2(x)$ can be computed in space*

$$O(s_1(\ell_2(n)) + s_2(n)),$$

where $\ell_2(n)$ is a bound on the length of the output of $f_2(x)$ on inputs of length n .

We can apply this lemma to the case of a single function being composed with itself many times:

Lemma A.2 ([10]). *Let $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be computable in space $s : \mathbb{N} \rightarrow \mathbb{N}$, where $s(n) \geq \log n$. Then, $g(x, k) = f \circ f \circ \dots \circ f(x)$ can be computed in space*

$$O\left(\sum_{i=0}^{k-1} s(\ell_i(n))\right),$$

where $\ell_i(n)$ is a bound on the length of the output of $g(x, i)$ on inputs of length n .

Samplers With Low Failure Probability. We prove that there exists a generator with our required properties. We first recall the statement:

THEOREM 2.3 ([8, 19, 22]). *Given $n, w, |\Sigma| \in \mathbb{N}$ and $\varepsilon, \delta > 0$, there exists a generator $\text{NIS} : \{0, 1\}^m \times \{0, 1\}^d \rightarrow \Sigma^n$ such that for every length n , width w branching program B we have:*

$$\Pr_{h \leftarrow U_m} \left[\left\| \overline{\text{NIS}}(h, \cdot) [B] - \overline{\mathbb{E}} [B] \right\| \leq \varepsilon \right] \geq 1 - \delta$$

and we have $m = O(\log(n) \log(nw/\varepsilon) + \log(|\Sigma|) + \log(1/\delta))$ and $d = O(\log(nw|\Sigma|/\varepsilon) + \log \log(1/\delta))$. Equivalently, for every branching program B , with probability $1 - \delta$ over $h \text{NIS}_h := \text{NIS}(h, \cdot)$ is ε -good for B . Furthermore, $\text{NIS}_h(x)$ can be evaluated in space $O(m)$ given two-way read-only access to the offline seed h .

We construct this in a standard fashion via the sampler trick. First, we recall both the PRG of Impagliazzo-Nisan-Wigderson [15], and a space-efficient averaging sampler. We use the INW PRG as it obtains optimal dependence on the size of the alphabet $|\Sigma|$.

THEOREM A.3 ([15]). *Given $n, w, |\Sigma| \in \mathbb{N}$ and $\varepsilon > 0$, there is a function $\text{GEN} : \{0, 1\}^s \rightarrow \Sigma^n$ with $m = O(\log n \log(nw/\varepsilon) + \log(|\Sigma|))$ that can be evaluated in space $O(s)$ such that for every ordered branching program $B : [w] \times \Sigma^n \rightarrow [w]$ of length n and width w ,*

$$\left\| \overline{\text{GEN}} [B] - \overline{\mathbb{E}} [B] \right\| \leq \varepsilon.$$

THEOREM A.4 ([8, 23]). *For every $\varepsilon, \delta > 0$ and $s \in \mathbb{N}$, there exists an averaging sampler $f : \{0, 1\}^m \times \{0, 1\}^d \rightarrow \{0, 1\}^s$ such that $d = O(\log(1/\varepsilon) + \log \log(1/\delta))$ and $m = s + O(\log(1/\delta) + \log(1/\varepsilon))$. Formally, for every function $g : \{0, 1\}^s \rightarrow \{0, 1\}$, we have*

$$\Pr_{h \leftarrow U_m} \left[\left| \mathbb{E}_{x \leftarrow U_d} [g(f(h, x))] - \mathbb{E}_{y \leftarrow U_s} [g(y)] \right| \leq \varepsilon \right] \geq 1 - \delta.$$

By combining these two ingredients we can construct the generator.

PROOF OF THEOREM 2.3. Let $\text{GEN} : \{0, 1\}^s \rightarrow \Sigma^n$ be the function of Theorem A.3 with $n = n, w = w, |\Sigma| = |\Sigma|$ and $\varepsilon = \varepsilon/2$. Observe that for every ordered branching program B of length n and width w , we have

$$\left\| \overline{\mathbb{E}} [B \circ \text{GEN}] - \overline{\mathbb{E}} [B] \right\| \leq \varepsilon/2.$$

Now let $f : \{0, 1\}^m \times \{0, 1\}^d \rightarrow \{0, 1\}^s$ be the function of Theorem A.4 with $\delta = \delta/w, \varepsilon = \varepsilon/2w$. For every function $B \circ \text{GEN}$ (where we choose a distinguished start and accept vertex and later take a union bound over ε and δ) we have

$$\Pr_{h \leftarrow U_m} \left[\left| \mathbb{E}_{x \leftarrow U_d} [B \circ \text{GEN}(f(h, x))] - \mathbb{E}_{y \leftarrow U_s} [B \circ \text{GEN}(y)] \right| \leq \varepsilon/2w \right] \geq 1 - \delta/w.$$

and thus

$$\Pr_{h \leftarrow U_m} \left[\left\| \overline{\text{GEN} \circ f(h, \cdot)} [B] - \overline{\mathbb{E}} [B] \right\| \leq \varepsilon/2 + \varepsilon/2 \right] \geq 1 - \delta$$

so letting $\text{NIS}(h, \cdot) = \text{GEN} \circ f(h, \cdot)$ we obtain the desired generator. \square

B BOUNDS AGAINST SINGLE TRANSITION BRANCHING PROGRAMS

To illustrate the added power obtained from derandomizing products $M_1 \cdots M_n$ of stochastic matrices $M_1, \dots, M_n \in \mathbb{R}^{w \times w}$ versus derandomizing powers M^n of a stochastic matrix $M \in \mathbb{R}^{w \times w}$, we observe that the latter model corresponds to derandomizing ordered branching programs with a single fixed transition function for every layer. We show via a short combinatorial argument that this limitation can be severe:

Lemma B.1. *For every even n , the function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ given by $f(x) = x_{n/2}$ cannot be computed by a single-transition ordered branching program of width $n/2 - 1$.*

We remark that f (and any parity on any subset of variables) can easily be computed by an ordered branching program (with distinct transition functions) of width 2.

PROOF. Let B be a single-transition branching program of width w computing f , and let v_0 be the start state and $A \subset [w]$ be the set of states marked as accept in the final layer.

Claim B.2. *For arbitrary $u, v \in [w]$, if there exists $\sigma \in \{0, 1\}^*$ such that $B[v, \sigma] = u$, there exists $\tau \in \{0, 1\}^{\leq w}$ such that $B[v, \tau] = u$.*

PROOF. This follows from the fact that if an $s - t$ path exists in a directed graph of size w , there must be an $s - t$ path of length at most w . \square

Now assume for contradiction that $w \leq n/2 - 1$. Let $u := B[v_0, 0^{n/2-1}1]$. We first claim that for any $\sigma \in \{0, 1\}^*$, $B[u, \sigma] \in A$. This follows from Claim B.2 and the fact that B must accept all strings of the form $B[v_0, 0^{n/2-1}1\sigma]$. But then there is $\tau \in \{0, 1\}^{\leq w}$ such that $B[v_0, \tau] = u$, and hence $B[v_0, \tau 0^{n-|\tau|}] \in A$, and as $n - |\tau| > n/2$ this is a contradiction to the fact that B computes f . \square

REFERENCES

- [1] AmirMahdi Ahmadinejad, Jonathan A. Kelner, Jack Murtagh, John Peebles, Aaron Sidford, and Salil P. Vadhan. 2020. High-precision Estimation of Random Walks in Small Space. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 1295–1306. <https://doi.org/10.1109/FOCS46700.2020.00123>
- [2] Roy Armoni. 1998. On the derandomization of space-bounded computations. In *Randomization and approximation techniques in computer science (Barcelona, 1998)*. Lecture Notes in Comput. Sci., Vol. 1518. Springer, Berlin, 47–59. https://doi.org/10.1007/3-540-49543-6_5
- [3] Mark Braverman, Gil Cohen, and Sumegha Garg. 2020. Pseudorandom Pseudodistributions with Near-Optimal Error for Read-Once Branching Programs. *SIAM J. Comput.* 49, 5 (2020). <https://doi.org/10.1137/18M1197734>
- [4] Mark Braverman, Sumegha Garg, and Or Zamir. 2021. Tight Space Complexity of the Coin Problem. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*. IEEE, 1068–1079. <https://doi.org/10.1109/FOCS52979.2021.00106>
- [5] Joshua Brody and Elad Verbin. 2010. The Coin Problem and Pseudorandomness for Branching Programs. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 30–39. <https://doi.org/10.1109/FOCS.2010.10>
- [6] Eshan Chattopadhyay, Pooya Hatami, Kaave Hosseini, and Shachar Lovett. 2019. Pseudorandom Generators from Polarizing Random Walks. *Theory Comput.* 15 (2019), 1–26. <https://doi.org/10.4086/toc.2019.v015a010>
- [7] Eshan Chattopadhyay, Pooya Hatami, Omer Reingold, and Avishay Tal. 2018. Improved pseudorandomness for unordered branching programs through local monotonicity. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, Ilias Diakonikolas, David Kempe, and Monika Henzinger (Eds.). ACM, 363–375. <https://doi.org/10.1145/3188745.3188800>
- [8] Eshan Chattopadhyay and Jyun-Jie Liao. 2020. Optimal Error Pseudodistributions for Read-Once Branching Programs. In *Proceedings of the 35th Computational Complexity Conference (CCC)*. 25:1–25:27. <https://doi.org/10.4230/LIPIcs.CCC.2020.25>
- [9] Gil Cohen, Dean Doron, Oren Renard, Ori Sberlo, and Amnon Ta-Shma. 2021. Error Reduction for Weighted PRGs Against Read Once Branching Programs. In *Proceedings of the 36th Computational Complexity Conference (CCC)*. 22:1–22:17. <https://doi.org/10.4230/LIPIcs.CCC.2021.22>
- [10] Gil Cohen, Dean Doron, and Ori Sberlo. 2022. Approximating Large Powers of Stochastic Matrices in Small Space. *Electron. Colloquium Comput. Complex.* TR22-008 (2022). ECCC:TR22-008 <https://ecc.ecc.weizmann.ac.il/report/2022/008>
- [11] Gil Cohen, Dean Doron, Ori Sberlo, and Amnon Ta-Shma. 2022. Approximating Iterated Multiplication of Stochastic Matrices in Small Space. *Electron. Colloquium Comput. Complex.* TR22-149 (2022). ECCC:TR22-149 <https://ecc.ecc.weizmann.ac.il/report/2022/149>
- [12] Michael A. Forbes and Zander Kelley. 2018. Pseudorandom Generators for Read-Once Branching Programs, in Any Order. In *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 946–955. <https://doi.org/10.1109/FOCS.2018.00093>
- [13] William M. Hoza. 2021. Better Pseudodistributions and Derandomization for Space-Bounded Computation. In *Proceedings of the 25th International Conference on Randomization and Computation (RANDOM)*. 28:1–28:23. <https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2021.28>
- [14] William M. Hoza and David Zuckerman. 2020. Simple Optimal Hitting Sets for Small-Success RL. *SIAM J. Comput.* 49, 4 (2020), 811–820. <https://doi.org/10.1137/19M1268707>
- [15] Russell Impagliazzo, Noam Nisan, and Avi Wigderson. 1994. Pseudorandomness for network algorithms. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, Frank Thomson Leighton and Michael T. Goodrich (Eds.). ACM, 356–364. <https://doi.org/10.1145/195058.195190>
- [16] Raghu Meka, Omer Reingold, and Avishay Tal. 2019. Pseudorandom generators for width-3 branching programs. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 626–637. <https://doi.org/10.1145/3313276.3316319>
- [17] Jack Murtagh, Omer Reingold, Aaron Sidford, and Salil Vadhan. 2019. Deterministic Approximation of Random Walks in Small Space. In *Proceedings of the 23rd International Conference on Randomization and Computation (RANDOM '19) (LIPIcs, Vol. 145)*, Dimitris Achlioptas and László A. Végh (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 42:1–42:22. <https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2019.42>
- [18] Jack Murtagh, Omer Reingold, Aaron Sidford, and Salil P. Vadhan. 2017. Derandomization Beyond Connectivity: Undirected Laplacian Systems in Nearly Logarithmic Space. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*. 801–812. <https://doi.org/10.1109/FOCS.2017.79>
- [19] Noam Nisan. 1992. Pseudorandom generators for space-bounded computation. *Combinatorica* 12, 4 (1992), 449–461. <https://doi.org/10.1145/100216.100242>
- [20] Noam Nisan and David Zuckerman. 1996. Randomness is Linear in Space. *J. Comput. System Sci.* 52, 1 (Feb. 1996), 43–52. <https://doi.org/10.1006/jcss.1996.0004>
- [21] Edward Pyne and Salil Vadhan. 2021. Pseudodistributions That Beat All Pseudorandom Generators (Extended Abstract). In *Proceedings of the 36th Annual Computational Complexity Conference (CCC)*. 33:1–33:15. <https://doi.org/10.4230/LIPIcs.CCC.2021.33>
- [22] Omer Reingold, Salil Vadhan, and Avi Wigderson. 2002. Entropy Waves, the Zig-Zag Graph Product, and New Constant-Degree Expanders. *Annals of Mathematics* 155, 1 (January 2002). <https://doi.org/10.1109/SFCS.2000.892006>
- [23] Omer Reingold, Salil Vadhan, and Avi Wigderson. 2004. A Note on Extracting Randomness from Santha–Vazirani Sources. (September 2004). Unpublished manuscript.
- [24] Michael Saks and Shiyu Zhou. 1999. $BSPACE(S) \subseteq DSPACE(S^{3/2})$. *J. Comput. System Sci.* 58, 2 (1999), 376–403. <https://doi.org/10.1006/jcss.1998.1616>
- [25] Walter J. Savitch. 1970. Relationships between nondeterministic and deterministic tape complexities. *J. Comput. System Sci.* 4 (1970), 177–192. [https://doi.org/10.1016/S0022-0000\(70\)80006-X](https://doi.org/10.1016/S0022-0000(70)80006-X)

Received 2022-11-07; accepted 2023-02-06