

Time-Space Tradeoffs for BPL via Catalytic Computation

Edward Pyne
MIT
epyne@mit.edu

November 13, 2023

Abstract

We prove that for every $\alpha \in [1, 1.5]$,

$$\mathbf{BSPACE}[S] \subseteq \mathbf{TISP} \left[2^{S^\alpha}, S^{3-\alpha} \right]$$

where $\mathbf{BSPACE}[S]$ corresponds to randomized space $O(S)$ computation, and $\mathbf{TISP}[T, S]$ to time $\text{poly}(T)$, space $O(S)$ computation. Our result smoothly interpolates between the results of (Nisan STOC 1992) and (Saks and Zhou FOCS 1995), which prove $\mathbf{BSPACE}[S]$ is contained in $\mathbf{TISP} [2^S, S^2]$ and $\mathbf{TISP} [2^{S^{3/2}}, S^{3/2}]$ respectively.

We obtain this result from new catalytic algorithms for derandomization. In the catalytic computation model introduced by (Buhrman, Cleve, Koucký, Loff, and Speelman STOC 2013), we are given a small worktape, and a larger catalytic tape that has an arbitrary initial configuration. We may edit this tape, but it must be exactly restored to its initial configuration at the completion of the computation. We prove that

$$\mathbf{BSPACE}[S] \subseteq \mathbf{CSPACE} [S, S^2]$$

where $\mathbf{CSPACE} [S, C]$ corresponds to catalytic algorithms that use $O(S)$ bits of workspace and $O(C)$ bits of catalytic space. Previously, only $\mathbf{BSPACE}[S] \subseteq \mathbf{CSPACE} [S, 2^{O(S)}]$ was known. In fact, we prove a general tradeoff, that for every $\alpha \in [1, 1.5]$,

$$\mathbf{BSPACE}[S] \subseteq \mathbf{CSPACE} [S^\alpha, S^{3-\alpha}].$$

We do not use the algebraic techniques of prior work on catalytic computation. Instead, we develop an algorithm that branches based on if the catalytic tape is conditionally random, and instantiate this primitive in a recursive framework.

1 Introduction

There has been extensive work over the last four decades seeking to understand the power of randomness in space-bounded computation. Let **BPL** be the set of languages recognized by randomized machines that run in space $O(\log n)$ on inputs of size n , and make two-sided error, and let **TISP** $[T(n), S(n)]$ be the set of languages recognized by deterministic machines that run in time $\text{poly}(T(n))$ and space $O(S(n))$ on inputs of size n . The 1969 result of Savitch [Sav70] implies

$$\mathbf{BPL} \subseteq \mathbf{TISP} \left[2^{\log^2 n}, \log^2 n \right].$$

In 1990, Nisan [Nis90] developed a pseudorandom generator (PRG) with seed length $O(\log^2 n)$ for read-once branching programs, which capture how **BPL** algorithms use randomness. By enumeration over seeds, this implies the same inclusion. Later, Nisan [Nis94] used his PRG in a nontrivial way to prove

$$\mathbf{BPL} \subseteq \mathbf{TISP} \left[n, \log^2 n \right],$$

i.e. $O(\log^2 n)$ space and polynomial time. In 1995, Saks and Zhou [SZ99] used the Nisan PRG in a very different way to prove

$$\mathbf{BPL} \subseteq \mathbf{TISP} \left[2^{\log^{3/2} n}, \log^{3/2} n \right],$$

and 26 years later, Hoza [Hoz21] improved this result to $\mathbf{BPL} \subseteq \mathbf{TISP} \left[2^{\log^{3/2-o(1)} n}, \log^{3/2-o(1)} n \right]$.

Despite extensive work on a variety of approaches for derandomizing space-bounded computation, such as PRGs [Nis90, INW94, NZ96, BRRY14, FK18, DMR⁺21, CHHL18], other pseudorandom objects such as weighted PRGs and hitting sets [BCG20, CL20, HZ18, PV21, CDR⁺21, Hoz21, CHL⁺23, CL23], non-black-box derandomizations [SZ99, PP23, CDST23], and results based on conjectured hardness [KvM02, DT23, PRZ23], the results of Nisan and Saks-Zhou and Hoza still constitute the frontier of knowledge for general derandomization of space-bounded computation. In fact, the frontier was “jagged”, as it was not known how to trade off between the incomparable results of [Nis94, SZ99].

1.1 Our Results

We smoothly interpolate between the results of Nisan and Saks and Zhou.

Theorem 1.1. *For every $\alpha \in [1, 1.5]$,*

$$\mathbf{BPL} \subseteq \mathbf{TISP} \left[2^{\log^\alpha n}, \log^{3-\alpha} n \right].$$

By a standard padding argument, this implies the result in the abstract for every space-constructible $S(n) \geq \log n$. For a concrete instantiation, we can set $\alpha = 1.25$ and obtain that **BPL** is contained in space $O(\log^{1.75}(n))$ and time $2^{O(\log^{1.25} n)}$. Previously, no result using this or less space ran in time less than $2^{\Omega(\log^{1.49} n)}$, and no result using this or less time ran in space less than $\Omega(\log^2 n)$.

Theorem 1.1 is an immediate consequence of new results in catalytic computation. In the catalytic logspace (**CL**) model, introduced by Buhrman, Cleve, Koucký, Loff, and Speelman [BCK⁺14], there is a machine M with $O(\log n)$ bits of standard working memory, and n^c bits of catalytic memory. This catalytic memory has an arbitrary initial configuration, and must be returned to exactly this configuration at the end of the computation. Perhaps surprisingly, [BCK⁺14] showed that **CL**

is likely to be strictly more powerful than **L**. In particular, it contains uniform **TC**¹ and thus **NL**. Motivated by this striking result, there have been several further works exploring the power of the model [BKLS18, GJST19, DGJ⁺20, CM20, BDS22].

We parameterize catalytic computation by time, space, and catalytic space (similar notions have been considered before, e.g. [BDS22]).

Definition 1.2. Let **CTISP** $[T(n), S(n), C(n)]$ be the set of languages recognized by catalytic machines that use $O(S(n))$ workspace and $O(C(n))$ catalytic space on inputs of size n , and run in time $\text{poly}(T(n))$ in the worst case.

Note that the worst-case runtime must hold over every catalytic tape, as well as every input. With this definition, we have

$$\mathbf{CTISP} [T, S, C] \subseteq \mathbf{TISP} [T, S + C]$$

as we can simulate the catalytic algorithm in the standard space-bounded model by initializing the catalytic tape to 0^C . We can then build algorithms that are efficient in time, space, *and* catalytic space. Our first result is to show that we can achieve a simulation of **BPL** that is as time- and space- efficient as that of Nisan [Nis94], and moreover almost all the space used is catalytic:

Theorem 1.3.

$$\mathbf{BPL} \subseteq \mathbf{CTISP} [n, \log n, \log^2 n].$$

This constitutes a superpolynomial improvement in the amount of catalytic space required to simulate **BPL**. Previously, it was known only that $\mathbf{BPL} \subseteq \mathbf{CTISP} [n, \log(n), \text{poly}(n)]$.¹

Our next step is to incorporate this algorithm into a recursive framework to derive a (time-efficient) tradeoff between the catalytic- and non-catalytic space consumption.

Theorem 1.4. *For every $\alpha \in [1, 1.5]$,*

$$\begin{aligned} \mathbf{BPL} &\subseteq \mathbf{CTISP} [2^{\log^\alpha(n)}, \log^\alpha(n), \log^{3-\alpha}(n)] \\ &\left(\subseteq \mathbf{TISP} [2^{\log^\alpha(n)}, \log^{3-\alpha}(n)] \right). \end{aligned}$$

Note that Theorem 1.4 immediately implies Theorem 1.1 (and in fact is stronger, as for all $\alpha > 1.5$ the substantial majority of the algorithm's space can be made catalytic). This result constitutes the first use of catalytic computation to prove new results about the standard model of space-bounded computation (answering in the affirmative Problem 5 of a recent survey of Mertz [Mer23]).

Interestingly, while previous work on algorithms for catalytic computation [BCK⁺14, CM20] primarily used algebraic techniques involving reversible computation over a ring, our results take a completely different approach. We hope that our techniques will have broader applications, both inside and beyond the model of catalytic computation.

1.2 Proof Overview for Theorem 1.3

A canonical (promise)-**BPL** complete problem is that of estimating transition probabilities in read-once branching programs:

¹We are aware of at least three proofs of this fact. The first follows from [BCK⁺14], an unpublished result (see the survey of [Mer23] for a sketch) proves it by treating the catalytic tape as a set of random walks, and the third follows from recent work on certified derandomization for **BPL** [PRZ23, DPT23].

Definition 1.5. A **read-once branching program (ROBP)** \overline{B} of width w and length n and alphabet $\{0, 1\}^t$ is defined by a function $B : [w] \times \{0, 1\}^t \rightarrow [w]$.² For $x \in (\{0, 1\}^t)^n$, define

$$\overline{B}[i, x] = B[B[\dots B[B[i, x_1], x_2] \dots], x_{n-1}, x_n].$$

It is well known that to derandomize **BPL**, it suffices to estimate $\Pr_{x \leftarrow U_n}[\overline{B}[1, x] = 1]$ up to error $1/3$ for an ROBP \overline{B} of length n , width n and alphabet $\{0, 1\}$.

The Result of Nisan. We now recall the result of [Nis94], which itself begins with the PRG of [Nis90]. In this PRG, we draw $\ell = \log n$ hash functions h_1, \dots, h_ℓ from a pairwise independent hash family on $t = O(\log nw)$ bits. We recursively define the PRG as follows. Let $\text{NIS}_0(x) = x$ for $x \in \{0, 1\}^t$, and let $\text{NIS}_{i+1}(x) = (\text{NIS}_i(x), \text{NIS}_i(h_{i+1}(x)))$. To analyze this PRG, fix a branching program $\overline{B} : (\{0, 1\}^t)^n \rightarrow \{0, 1\}$ of width w , with transition function B . Viewing this construction from the bottom up, the first hash function h_1 is good if for every $a, b \in [w]$,

$$\Pr_{x, x' \leftarrow U_t} [B[B[a, x], x'] = b] \approx \Pr_{x \leftarrow U_t} [B[B[a, x], h_1(x)] = b],$$

i.e. the distribution $(x, h_1(x))$ is indistinguishable from the distribution (x, x') by the composition of B with itself. Since B can only pass $\log w$ bits of information from the first to the second half, this occurs with probability $1 - w^{-c}$ over $h_1 \leftarrow \mathcal{H}$ (assuming $t = O(\log w)$ is sufficiently large). The ultimate PRG is analyzed recursively using ℓ applications of essentially the same idea. Concretely, at the second level of the construction, we now want a hash function h_2 that fools the length $n/2$ program with transition function $B'[a, x] = B[B[a, x], h_1(x)]$.

While the Nisan PRG randomly selects ℓ hash functions at once, the insight of [Nis94] was that, given a specific program \overline{B} that we want to fool, we can *search* for good hash functions level by level. At level i , we find a hash function h_i that fools the relevant transition function. As this test is easy to implement in time $2^{O(t)}$ for a fixed h and there are $2^{O(t)}$ such h to test, we can find such a good hash function in time $2^{O(t)} = \text{poly}(n)$ per level, giving a polynomial runtime overall.

An Algorithm From Conditional Compression. We transform this algorithm into a *catalytic* algorithm as follows. Suppose we have a branching program \overline{B} of width w and length $n = 2^\ell$, and a catalytic tape \mathbf{w} , with an arbitrary initial configuration. We interpret \mathbf{w} as holding 2ℓ hash functions $h_1, \dots, h_{2\ell}$, each over $t = O(\log nw)$ bits (and note that each function can have description size exactly $2t$). Let $V \in \{0, 1, *\}^{2\ell}$ and initialize $V = *^{2\ell}$ to indicate the status of each block. We then iterate through this list. Letting the i th hash function be \tilde{h} and the previous good hash functions be \vec{h}_p , we check if \tilde{h} is a good hash function, using the test as before.

- If \tilde{h} is good, we set $V_i = 1$, indicating \tilde{h} is part of the list of good hashes.
- If \tilde{h} is not good, it must lie in the set $\text{FAIL}(\vec{h}_p)$ of hash functions that fail to fool the current transition function. But as almost all h are good, the index of \tilde{h} in $\text{FAIL}(\vec{h}_p)$ is a concise description of \tilde{h} ! We can then replace \tilde{h} with this index, and free up $\Omega(\log nw)$ bits on this block of the tape. Finally, set $V_i = 0$ to indicate we have compressed this block.

At the end of this phase, we have either found ℓ good hash functions, or have freed up $\ell \cdot \Omega(\log nw)$ bits on the tape. In the latter case, we can simply search for a good set of hash functions (on

²The standard definition of ROBPs permits the transition function to differ between the layers. However, as we will always be dealing with programs where $w \geq n$, and we are insensitive to polynomial losses in the width, we can assume all transition functions are the same for clarity.

slightly fewer bits), exactly as in the algorithm of [Nis94], and store these in the free space of the compressed blocks. Thus, in both cases we obtain a sequence of hash functions that together constitute a good PRG for \overline{B} , and hence can construct a generator NIS that does a good job estimating walk probabilities on \overline{B} . The final step of estimating these walks can be performed in space $O(t + \log nw) = O(\log nw)$ with read-only access to the tape \mathbf{w} .

Finally, to return the tape to its original configuration, we work backwards over the compressed blocks, i.e. indices i where $V_i = 0$. For each block, we determine the preceding good hash functions \vec{h}_p , read the index of the original hash (i.e. tape configuration) in $\text{FAIL}(\vec{h}_p)$, then find the hash with this index by enumeration and write it to the tape.

1.3 Proof Overview for Theorem 1.4

It is not immediately obvious why a catalytic analogue of Nisan’s decades old result would help in obtaining new results in derandomization, as ultimately the total space consumed by the algorithm of Theorem 1.3 is the same as that of [Nis94]. Our next idea is to unify this with efficient *composition* of catalytic algorithms:

Composition of Catalytic Algorithms. Recall that in the conventional composition of space-bounded algorithms, we can compute the composition of two algorithms running in space $S(n)$ in space $c \cdot S(n)$, for some constant $c > 1$. Our key observation is that for catalytic algorithms, we can obtain composition with *no* increase in the length of the catalytic tape:

Theorem 1.6 (Composition of Catalytic Space-Bounded Algorithms). *Given two catalytic algorithms $\mathcal{M}_1, \mathcal{M}_2$ computing f_1, f_2 respectively, each using space $S(n) \geq \log n$, catalytic space $C(n)$, and time $T(n)$, there is a catalytic algorithm \mathcal{M}' using time $\text{poly}(T(n))$, space $O(S(n))$, and catalytic space $C(n)$ that computes $f_2 \circ f_1$.*

The proof of this result modifies the standard composition of space-bounded algorithms. To compute $\mathcal{M}_2(\mathcal{M}_1(x))$, we begin to simulate $\mathcal{M}_2^{\mathbf{w}}(f_1(x))$ (where the superscript notation denotes running the machine with catalytic tape \mathbf{w}). Whenever \mathcal{M}_2 reads a bit of the input, we simulate $\mathcal{M}_1^{\mathbf{w}'}(x)$ to obtain the relevant bit of $f_1(x)$, where \mathbf{w}' is the current configuration of the catalytic tape of \mathcal{M}_2 . Since \mathcal{M}_1 is guaranteed to produce the correct answer for every starting tape, we have that $\mathcal{M}_1^{\mathbf{w}'}(x) = f_1(x)$. Moreover, as \mathcal{M}_1 is catalytic, it resets the tape to \mathbf{w}' before returning, so \mathcal{M}_2 does not notice the call has occurred, and can continue its computation.

We remark that we are not able to apply this theorem as-is due to issues with (essentially) f_1 being a relation with multiple valid outputs, so the actual statement we prove is more involved. In particular, we must deal with safety reverting the catalytic tape if an intermediate call to \mathcal{M}_1 fails.

Derandomization via Repeated Powering. Going from this to an improved algorithm for derandomization requires a final further ingredient, which is given by a variant of the Saks-Zhou recursive powering scheme. Saks-Zhou [SZ99] divides computing the n th power of an $n \times n$ stochastic matrix M (a **prBPL** complete problem) into r_2 iterations of computing the 2^{r_1} th power, for any $r_1 r_2 = \log n$. For convenience, let $M_0 = M$ and $M_i = M^{2^{r_1 \cdot i}}$ for $i \in [r_2]$. In the original algorithm, all levels share a single set of hash functions $\vec{h} = (h_1, \dots, h_{r_1})$, each on $O(\log n)$ bits. A random set of hash functions will do a good job computing M_i from M_{i-1} for every i , and so we can reuse this fixed set of hash functions at every level.³ Unfortunately, such an argument is incompatible

³There are additional complications from reusing the hash functions, but they are not the primary reason for the high time complexity.

with searching for good hash functions one by one. Since we use every hash function to produce an approximation to M_1 , if we later discover a hash function is bad at powering M_i for $i \geq 1$, seemingly we must destroy all partial progress and try a new set of hash functions. Thus, the Saks-Zhou algorithm must enumerate over $\vec{h} = (h_1, \dots, h_{r_1})$ all at once, incurring a runtime of $2^{\Omega(r_1 \cdot \log n)}$. As the algorithm incurs a runtime of $2^{\Omega(r_2 \cdot \log n)}$ merely from the recursive composition of space bounded algorithms, the total runtime is at least $2^{\Omega(\max\{r_1, r_2\} \cdot \log n)} = 2^{\Omega(\log^{3/2} n)}$ for any setting of parameters r_1 and r_2 .

Composing Conditional Compression Algorithms. Our catalytic algorithm allows for a more efficient approach. We follow the same recursive powering scheme as Saks-Zhou, but at each level use the algorithm of Theorem 1.3 that treats \mathbf{w} as a list of $2r_1$ candidate hash functions.⁴ Whenever we request an entry of a smaller power, we call the next level algorithm. If that level sees that the hash functions currently on the tape are good, it uses them to compute the requested entry. If not, it temporarily compresses the tape, finds good hash functions in time $\text{poly}(n)$, uses them to compute the requested entry, then resets the tape to exactly the same configuration the calling algorithm was expecting before returning. Thus, every level can either use the tape as-is, if it is suitable, or quickly compute a better set of hash functions on the fly and revert before returning control. This eliminates the $2^{r_1 \log n}$ term in the runtime. Moreover, the $O(r_1 \log n)$ bits used to store the hash functions can be treated as catalytic space, resulting in an algorithm that uses only $O(r_2 \log n)$ bits of workspace.

Finally, for every $\alpha \in [1, 1.5]$, we can choose $r_1 = \log^{2-\alpha}(n)$ and $r_2 = \log^{\alpha-1}(n)$ and obtain a algorithm that uses $O(r_1 \log n) = O(\log^{3-\alpha} n)$ catalytic space, $O(r_2 \log n) = O(\log^\alpha n)$ workspace, and runs in time $\text{poly}(n^{r_2}) = 2^{O(\log^\alpha n)}$, as claimed.

Such an approach runs into a subtle technical issue. Since the algorithm at level i may be called many times with different starting catalytic tapes, we must ensure that the algorithm returns the *same* approximate power each time, as otherwise the composition would not be well defined. To fix this, we first define a notion of catalytic algorithms that are allowed to return \perp for some initial catalytic tapes, in addition to a fixed output that is independent of the catalytic tape. We then show how these algorithms can be composed, while still maintaining the ability to revert the tape to the original configuration in the worst case. Finally, we adopt the strategy of Saks and Zhou [SZ99], and randomly perturb the matrices at each level. In our case, if a level of the algorithm determines that a shift is bad (i.e. could produce ambiguous behavior) it aborts and returns \perp . We show with high probability over the shifts, this will never occur (i.e. we will not return \perp) no matter the tape, and so we can compose the algorithm with itself and find the desired output.

Showing that we can successfully avoid permanent damage to the tape in the case that the shifts are bad requires further work. In particular, we ensure that our catalytic algorithms can be *reverted* from any point, where our notion of reversibility requires that we do not introduce any new configurations of the catalytic tape. We show that we can achieve this notion without a substantial time cost, and moreover it is compatible with recursive composition. Using this tool, we are able to return to the original tape configuration of a subroutine ever returns \perp .

1.4 Roadmap

In Section 2, we formally define the catalytic computation model, and prove Theorem 1.6. In Section 3, we prove Theorem 1.3, and in Section 4, we prove Theorem 1.1. In Appendix A we provide proofs of some cited lemmas.

⁴We give a “non-black-box” explanation of the final algorithm here as it illustrates the actual idea, but our proof uses a black-box statement regarding composition of catalytic algorithms.

2 Catalytic Machines and Composition

We first formally define a catalytic Turing machine.

Definition 2.1 (Catalytic Turing Machine [BCK⁺14]). A Turing machine \mathcal{M} is a **catalytic machine** using time $T(n)$, workspace $S(n)$, and catalytic space $C(n)$ if it has a work tape, a read-only input tape, a write-only output tape, and a catalytic tape \mathbf{w} . We require that for every input x with $|x| = n$ and every \mathbf{w} , $\mathcal{M}^{\mathbf{w}}(x)$ halts in time at most $T(n)$, using at most $S(n)$ cells on the worktape and $C(n)$ cells on \mathbf{w} . Moreover, the final configuration of \mathbf{w} must be equal to its initial configuration, for every x and \mathbf{w} .

We now define the notion of a catalytic machine that computes a function. We furthermore define the notion of partially computing a function, where on some tapes \mathbf{w} the machine can output a special failure symbol \perp .

Definition 2.2. For a function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$, we say a catalytic machine \mathcal{M} (**catalytically computes** f) if for every x and \mathbf{w} , $\mathcal{M}^{\mathbf{w}}(x) = f(x)$, and at the end of the computation \mathbf{w} is in its original state. We say that \mathcal{M} **partially (catalytically) computes** f if for every x and \mathbf{w} , $\mathcal{M}^{\mathbf{w}}(x) \in \{\perp, f(x)\}$, and at the end of the computation (no matter the output) \mathbf{w} is in its original state.

Partial catalytic computation is trivial without further restrictions (as \mathcal{M} can always output \perp), but we require it as an intermediate step in our analyses. We require a further condition on our machines, that they can revert the catalytic tape at any time without the catalytic tape traversing any new configurations:⁵

Definition 2.3. A catalytic machine \mathcal{M} is **reversible** if for every x and initial configuration \mathbf{w} , at any point during the execution of $\mathcal{M}^{\mathbf{w}}(x)$, the machine can receive an external REVERT signal. Let $P = P(\mathbf{w})$ denote all prior configurations of the catalytic tape during the execution of $\mathcal{M}^{\mathbf{w}}(x)$. After this signal, \mathcal{M} must reset \mathbf{w} to the original configuration, and moreover every intermediate configuration of \mathbf{w} during this process must lie in P . We require any time bound on \mathcal{M} to hold even in the case that \mathcal{M} is given the REVERT command at an arbitrary point.

2.1 Composition of Catalytic Algorithms

We state the main result of this section, which is that catalytic algorithms can be composed without increasing the *catalytic* space usage. We must be careful when dealing with partial catalytic machines, and in this case we only obtain composition if the machines are reversible (Definition 2.3).

Theorem 2.4 (Composition of Partial Catalytic Machines). *Suppose reversible catalytic machines $\mathcal{M}_1, \mathcal{M}_2$ partially compute $f_1, f_2 : \{0, 1\}^n \rightarrow \{0, 1\}^n$ respectively using workspace $S(n) \geq \log(n)$, catalytic space $C(n)$, and time $T(n)$. Then there is a reversible catalytic machine \mathcal{M} that partially computes $f_2 \circ f_1$ using workspace $O(S)$, catalytic space C , and time $\text{poly}(T(n))$. Moreover, $\mathcal{M}^{\mathbf{w}}(x) = \perp$ only if $\mathcal{M}_2^{\mathbf{w}}(f_1(x)) = \perp$, or there exists \mathbf{w}' such that $\mathcal{M}_1^{\mathbf{w}'}(x) = \perp$.*

Proof. We proceed roughly following the standard proof for composition of space-bounded algorithms. We maintain two sections on the worktape of size S for \mathcal{M}_1 and \mathcal{M}_2 , and a single catalytic tape \mathbf{w} .

⁵There are existing results related to transforming catalytic algorithms into reversible catalytic algorithms [Mer23]. However, they do not appear to maintain worst-case runtime over the catalytic tape, which is crucial for our results.

The Simulation. We now begin to simulate $\mathcal{M}_2^{\mathbf{w}}$. We first verify that $\mathcal{M}_1^{\mathbf{w}}(x) \neq \perp$. As in the conventional composition of space-bounded algorithms, every time \mathcal{M}_2 reads its input, we run $\mathcal{M}_1^{\mathbf{w}}(x)$ on a separate section of the worktape and return the relevant bit of its output, where \mathbf{w} is the same catalytic tape used by \mathcal{M}_2 , in whatever its current configuration is at the time of the tape read. Moreover, every time \mathcal{M}_2 writes to the catalytic tape, resulting in a configuration \mathbf{w}' , we run $\mathcal{M}_1^{\mathbf{w}'}(x)$ and verify that it does not produce \perp .

Computing the Function. In the case that $\mathcal{M}_1^{\mathbf{w}'}(x) = f_1(x)$ for every configuration \mathbf{w}' that is encountered in this simulation and $\mathcal{M}_2^{\mathbf{w}}(f_1(x)) = f_2(f_1(x))$, it is easy to see that $\mathcal{M}^{\mathbf{w}}(x) = f_2(f_1(x))$. Moreover, it is clear that in this case we successfully reset the tape. Otherwise, consider the first point at which $\mathcal{M}_1^{\mathbf{w}'}(x) = \perp$. We first undo the most recent change to the catalytic tape, and send the REVERT command to \mathcal{M}_2 . Once \mathcal{M}_2 has finished reverting, return \perp . We claim that \mathcal{M}_2 successfully reverts the tape. This follows from the reversibility of \mathcal{M}_2 , and the fact that all calls \mathcal{M}_2 makes to its input during this process are correctly answered by \mathcal{M}_1 . The latter property follows as every time \mathcal{M}_2 queries its input during the revert process, \mathbf{w} is in a state that was encountered during the forward pass, and hence $\mathcal{M}_1(x)$ produced $f_1(x)$ when initialized with this catalytic configuration (as otherwise we would have aborted sooner).

Reversibility. Essentially the same argument establishes that \mathcal{M} is reversible. If we receive the REVERT command, let $P(\mathbf{w})$ be the states of the catalytic tape that have been encountered so far. First send REVERT to \mathcal{M}_1 (if operating), and once it has completed send the REVERT command to \mathcal{M}_2 . Moreover, while \mathcal{M}_2 is reverting, we claim that \mathbf{w} remains in $P(\mathbf{w})$. This follows from the fact that \mathcal{M}_2 is reversible (as any configurations it creates will lie in $P(\mathbf{w})$), and moreover every time \mathcal{M}_2 queries its input, any computation done by \mathcal{M}_1 will likewise keep \mathbf{w} in $P(\mathbf{w})$, as we already called \mathcal{M}_1 with this starting configuration in the forward pass (and \mathcal{M}_1 will not produce \perp , as otherwise we would have already aborted). Thus, the composed algorithm is reversible.

Time and Space. We now argue the space and time are as claimed. There are a constant number of pointers (which we maintain on the worktape) to track the number of bits output by \mathcal{M}_1 , current tape heads, and other information. The fact that the catalytic tape size is preserved is immediate. The call overhead adds at most a polynomial factor in the runtime, as we run \mathcal{M}_1 at most once per step of \mathcal{M}_2 . Finally, if \mathcal{M}_1 computes $f_1(x)$ for every catalytic tape and \mathcal{M}_2 computes correctly on \mathbf{w} , we successfully compute $f_2 \circ f_1$ as claimed. \square

We derive an easy corollary in the case of multiple composition:

Corollary 2.5. *Suppose a reversible catalytic machine \mathcal{M} partially computes $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ using workspace $S(n) \geq \log n$, catalytic space C , and time 2^S . Then there exists a reversible catalytic machine \mathcal{M}' that partially computes f^ℓ using workspace $O(\ell \cdot S)$, catalytic space C , and time $2^{O(\ell \cdot S)}$. Moreover, for x where $\mathcal{M}^{\mathbf{w}}(f^i(x)) = f(f^i(x))$ for every \mathbf{w} and $i \in \{0, \dots, \ell - 1\}$, $\mathcal{M}'^{\mathbf{w}}(x) = f^\ell(x)$.*

3 Catalytic Derandomization From Conditional Compression

In this section we prove Theorem 1.3. We state all our results in terms of catalytic algorithms for the stochastic matrix powering problem, as it is easily compatible with the recursive framework we implement later. Recall a nonnegative matrix is **stochastic** (resp. **substochastic**) if all row sums

are 1 (resp. at most 1). For a set S , let U_S be the uniform distribution over S , and let $U_n = U_{\{0,1\}^n}$.

Theorem 3.1. *There is a **CTISP** $[n, \log n, \log^2 n]$ algorithm that, given n and a stochastic matrix $M \in [0, 1]^{n \times n}$ where each entry is specified with $O(\log n)$ bits of precision, outputs $\widetilde{M} \in [0, 1]^{w \times w}$ such that $\left\| \widetilde{M} - M^n \right\|_1 \leq 1/n$.*

We recall the existence of efficient algorithms which canonicalize (sub)stochastic matrices, essentially reducing the stochastic matrix powering problem to producing a PRG that fools a branching program.

Lemma 3.2 ([SZ99, PP23, CDST23]). *There is a constant $c > 0$ and a space $O(\log nw/\varepsilon)$ algorithm which, given $\varepsilon > 0$ and $n, w \in \mathbb{N}$ where $w \geq n$ and a substochastic matrix $M \in [0, 1]^{w \times w}$ with $O(\log w)$ bits of precision, returns a branching program*

$$\overline{B} : [(w/\varepsilon)^c] \times \{0, 1\}^m \rightarrow [(w/\varepsilon)^c]$$

where $m = n \cdot O(\log(w/\varepsilon))$ is a power of two satisfying the following. Letting $\widetilde{M} \in [0, 1]^{w \times w}$ be the (substochastic) matrix where for $i, j \in [w]$ we define $\widetilde{M}_{i,j} = \Pr_{x \leftarrow U_n}[\overline{B}[i, x] = j]$, we have $\left\| \widetilde{M} - M^n \right\|_1 \leq \varepsilon$.

As this is not the way these results are stated, we provide a translation in Appendix A. We next define the Nisan PRG, and recall several auxiliary lemmas.

The Nisan PRG. Given a branching program, we first define the larger alphabet program obtained from duplicating each edge:

Definition 3.3. For $t \in \mathbb{N}$, for $\overline{B} : [w] \times \{0, 1\}^n \rightarrow [w]$ of width w , let $\overline{B}_t : [w] \times (\{0, 1\}^t)^n \rightarrow [w]$ be the branching program of length n and width w over alphabet $\{0, 1\}^t$ with transition function $B_t[a, y] = \overline{B}[a, y_1]$. Note that \overline{B}_t can be constructed in space $O(\log tnw)$ given \overline{B} , and furthermore for every $i, j \in [w]$, $\Pr_{x \leftarrow U_n}[\overline{B}[i, x] = j] = \Pr_{x \leftarrow U_{(\{0,1\}^t)^n}}[\overline{B}_t[i, x] = j]$.

We recall a pairwise independent hash family with a very efficient description:

Fact 3.4. For every $t \in \mathbb{N}$, there exists a pairwise independent hash family $\mathcal{H} : \{0, 1\}^t \rightarrow \{0, 1\}^t$ such that $|\mathcal{H}| = 2^{2t}$, and $h \in \mathcal{H}$ (which we associate with $h \in \{0, 1\}^{2t}$) can be evaluated in space $O(t)$.

Given a (hash) function $h : \{0, 1\}^t \rightarrow \{0, 1\}^t$ and a program \overline{B} , we define an operator that applies a single level of the Nisan construction with hash function h .

Definition 3.5. Given $\overline{B}_t : [w] \times (\{0, 1\}^t)^n \rightarrow [w]$ and $h : \{0, 1\}^t \rightarrow \{0, 1\}^t$, let $\overline{B}_{t,h} : [w] \times (\{0, 1\}^t)^{n/2} \rightarrow [w]$ be the width w , length $n/2$ program with transition function

$$B_{t,h}[a, x] = \overline{B}_t[B_t[a, x], h(x)].$$

Using a recursive application of hash functions, we can define the Nisan PRG as follows.

Definition 3.6. For $(h_1, \dots, h_\ell) \in \mathcal{H}_t$, define $\text{NIS}_{(h_1, \dots, h_\ell)} : \{0, 1\}^t \rightarrow \{0, 1\}^{t \cdot n}$ inductively as follows. Let $\text{NIS}_0(x) = x_1$, and for $j \in [\ell]$

$$\text{NIS}_{(h_1, \dots, h_j)}(x) = (\text{NIS}_{(h_1, \dots, h_{j-1})}(x) \| \text{NIS}_{(h_1, \dots, h_{j-1})}(h_j(x))).$$

Note that $B[\cdot, \text{NIS}_{(h_1, \dots, h_\ell)}(\cdot)]$ and $B_{t, (h_1, \dots, h_\ell)}[\cdot, \cdot]$ (as defined in Definition 3.5) are equal as functions.

To analyze the Nisan PRG, we define the notion of a hash function being good for composing two functions, and a PRG being good for a function.

Definition 3.7. For every $n, w, t \in \mathbb{N}$ and $\delta > 0$ and $\bar{f} : [w] \times (\{0, 1\}^t)^n \rightarrow [w]$ and $G : \{0, 1\}^t \rightarrow (\{0, 1\}^t)^n$, we say that G is δ -**good** for \bar{f} if for every $i, j \in [w]$,

$$\left| \Pr_{x \leftarrow U_t} [\bar{f}[i, G(x)] = j] - \Pr_{x \leftarrow U_{(\{0, 1\}^t)^n}} [\bar{f}[i, x] = j] \right| \leq \delta.$$

Moreover, we say $h \in \mathcal{H}$ is δ -**good** (and δ -**bad** otherwise) for $f : [w] \times \{0, 1\}^t \rightarrow [w]$ if $G(x) = (x || h(x))$ is δ -good for $\bar{f}[i, (x, y)] = f[f[i, x], y]$.

We recall that a random hash function is good with high probability.

Lemma 3.8 ([Nis90]). *For every f , $\Pr_{h \leftarrow \mathcal{H}_t}[h \text{ is } \delta\text{-good for } f] \geq 1 - w^5(1/\delta)^2/2^t$.*

(We provide a proof in Appendix A.) Moreover, a hybrid argument establishes the following.

Lemma 3.9 ([Nis90]). *For every $\bar{B}_t : [w] \times (\{0, 1\}^t)^n \rightarrow [w]$ and $\vec{h} = (h_1, \dots, h_\ell)$, suppose for every $i \in [\ell]$, h_i is δ -good for $B_{t, h_1, \dots, h_{i-1}}$. Then $\text{NIS}_{\vec{h}}$ is $\delta \cdot nw$ -good for \bar{B} .*

Catalytic Derandomization. We now state the main result that powers all of our derandomizations.

Theorem 3.10. *There is a pair of reversible catalytic algorithms \mathcal{A}, \mathcal{D} that run in workspace $O(\log nw/\varepsilon)$, catalytic space $O(\log(n) \cdot \log(nw/\varepsilon))$, and time $\text{poly}(nw/\varepsilon)$ and act as follows. Given $\varepsilon > 0$ and a length $n = 2^\ell$, with w ROBP $\bar{B} : [w] \times \{0, 1\}^n \rightarrow [w]$ where $w \geq n$:*

- *The machine $\mathcal{A}^{\mathbf{w}}(\bar{B})$ outputs $V \in \{0, 1\}^{2^\ell}$ and $t = O(\log nw/\varepsilon)$ and sets the catalytic tape to \mathbf{w}' , such that (\mathbf{w}', V) contains a (read-only) data structure supporting access to hash functions $\vec{h} = (h_1, \dots, h_\ell)$ each on t bits, such that $\text{NIS}_{\vec{h}}$ is ε -good for \bar{B} .*
- *The machine $\mathcal{D}^{\mathbf{w}'}(\bar{B}, V)$ sets the final catalytic tape configuration to \mathbf{w} .*

To make our compression and decompression algorithms work, we require that we can determine if a hash function is good for a branching program at a certain level of the Nisan construction, given pointers to the hash functions at the previous levels:

Proposition 3.11. *There is a space $O(t + \log(w/\delta))$ algorithm that, given $n, w, t \in \mathbb{N}$ with $w \geq n$ and $\vec{h} \in \{0, 1\}^{2t}$ and $\bar{B} : [w] \times \{0, 1\}^n \rightarrow [w]$ and (read only) \mathbf{w} and pointers p_1, \dots, p_r such that $\mathbf{w}_{[p_i, \dots, p_i+2t]} = h_i$ represents a hash function on t bits, returns if \vec{h} is δ -good for $B_{t, (h_1, \dots, h_r)}$.*

We give a proof in Appendix A, as it essentially follows from the argument of [Nis94]. We can then prove the theorem:

Proof of Theorem 3.10. We assume without loss of generality that n, w and $1/\varepsilon$ are powers of two. Set

$$t_0 = 60 \log(w/\varepsilon), \quad t_1 = 25 \log(w/\varepsilon), \quad \delta = \varepsilon/nw \geq \varepsilon/w^2,$$

and note that we choose t_1 large enough such that a good series of hash functions on t_1 bits always exists. The algorithm works as follows. First, virtually divide the catalytic tape as:

$$\mathbf{w} = \left(\mathbf{w}^1 || \mathbf{w}^2 || \dots || \mathbf{w}^{2^\ell} \right)$$

where $|\mathbf{w}^i| = 2t_0$, which we think of as initially holding $h : \{0, 1\}^{t_0} \rightarrow \{0, 1\}^{t_0}$. Note that $|\mathbf{w}| = \ell \cdot 4t_0 = O(\log(n) \log(nw/\varepsilon))$ as claimed.

Next, initialize $V \in \{0, 1, *\}^{2\ell}$ to indicate if each block is compressed, uncompressed, or unprocessed respectively. The first two cases correspond to the following two formats of the block:

$$\mathbf{w}^i = \begin{cases} h & V_i = 1 \\ (z || 0^{50 \log(w/\varepsilon)}) & V_i = 0 \end{cases}$$

Informally, the first corresponds to block i originally containing a good hash function for \overline{B} , and the second corresponds to block i originally containing a bad hash function, which is thus compressible (in fact, z represents a compressed version of the original data). We define notation for the set of blocks in each configuration:

Definition 3.12. For $b \in \{0, 1\}$, let $I^b(V) \subseteq [2\ell]$ correspond to the indices such that $V_i = b$, and let $S^b(V) = |I^b(V)|$.

Next, we initialize a counter $i = 1$ for the current block. We then iterate over $i = \{1, \dots, 2\ell\}$ until $\max\{S^1(V), S^0(V)\} = \ell$.⁶ For each i , the algorithm works as follows. Let $\tilde{h} = \mathbf{w}^i$ be the hash function (on t_0 bits) obtained from the current block. We then test if \tilde{h} is δ -good for

$$f = B_{t_0, \vec{h}_p}, \text{ where } \vec{h}_p = \left(\mathbf{w}^{I^1(V)_1}, \dots, \mathbf{w}^{I^1(V)_{S^1(V)}} \right)$$

corresponds to the hash functions on the preceding good blocks, and B_{t_0, \vec{h}_p} is defined as in Definition 3.5. As the index set $I^1(V)$ is easy to generate given V , this test can be performed in space $O(\log nw/\varepsilon)$ without modifying the catalytic tape (and hence also in time $\text{poly}(nw/\varepsilon)$), by Proposition 3.11.

Given the results of this test, we break into cases depending on if \tilde{h} is good:

- If \tilde{h} is δ -good for f , set $V_i = 1$.
- If \tilde{h} is δ -bad for f , set $V_i = 0$. Next, by enumeration over strings $h \in \{0, 1\}^{2t_0}$ (which we can do using the workspace), determine the index of \tilde{h} in the set

$$\mathbf{BAD}_i = \left\{ h \in \{0, 1\}^{2t_0} : h \text{ is } \delta\text{-bad for } B_{t_0, \vec{h}_p} \right\}$$

where we again perform this test using Proposition 3.11. Letting the index of \tilde{h} in this set be z , write

$$\mathbf{w}^i = \left(z || 0^{50 \log(w/\varepsilon)} \right).$$

We denote the final $50 \log(w/\varepsilon)$ bits as free space.

Finally, we claim that we can in fact write these quantities in space $|\mathbf{w}^i| = 2t_0$. We have

$$\begin{aligned} |\mathbf{BAD}_i| &= 2^{2t_0} \cdot \Pr_{h \leftarrow \mathcal{H}} [h \text{ is } \delta\text{-bad for } f] \\ &\leq 2^{2t_0} \cdot w^5 (1/\delta)^2 / 2^{t_0} && \text{(Lemma 3.8)} \\ &\leq 2^{2t_0} \cdot (w/\varepsilon)^{7-60} \end{aligned}$$

And thus $\log |\mathbf{BAD}_i| \leq 2t_0 - 50 \log(w/\varepsilon) = |\mathbf{w}^i| - 50 \log(w/\varepsilon)$. Therefore, we can record all required information as claimed.

⁶If we exit before $i = 2\ell$, set the remaining indices of V to an arbitrary value, which we ignore for clarity of presentation.

After processing all blocks, we obtain a catalytic tape \mathbf{w}' and one of two cases:

- If $S^1(V) = \ell$, there exist $\vec{h} = (h_1, \dots, h_\ell)$ corresponding to the hash functions (on t_0 bits) in $I^1(V)$, and these functions are easy to recover from V .
- Else, we must have $S^0(V) = \ell$.

For $i \in I^0(V)$, let F_i be the $50 \log(w/\varepsilon)$ free bits in \mathbf{w}^i . Note that a description of a hash function $h : \{0, 1\}^{t_1} \rightarrow \{0, 1\}^{t_1}$ is of size $|F_i|$. Iterating over $i \in I^0(V)$ in increasing order, we find (via brute force enumeration) a hash function \tilde{h} that is δ -good for

$$f = B_{t_1, \vec{h}_p}, \text{ where } \vec{h}_p = \left(\mathbf{w}_{F_{I^0(V)_1}}, \dots, \mathbf{w}_{F_{I^0(V)_{i-1}}} \right).$$

corresponds to the (δ -good) hash functions stored on the free space in the preceding indices of $I^0(V)$. Next, store \tilde{h} in \mathbf{w}_{F_i} . Such a good hash function always exists, by our choice of t_1 and Lemma 3.8, and moreover testing if each candidate is good can be computed in the desired space and time by Proposition 3.11. After this processing,

$$\left(\mathbf{w}_{F_{I^0(V)_1}}, \dots, \mathbf{w}_{F_{I^0(V)_\ell}} \right)$$

contain ℓ good hash functions, which we can clearly access in read-only fashion given V and \mathbf{w}' .

Thus, in both cases we obtain a set of hash functions $\vec{h} = (h_1, \dots, h_\ell)$ on $t = O(\log nw/\varepsilon)$ bits that is δ -good for every one of the relevant tests, so by Lemma 3.9 we have that $\text{NIS}_{\vec{h}}$ is $\delta \cdot nw \leq \varepsilon$ -good for \bar{B} .

Decompression and Reversibility. It suffices to show that at any point, the algorithm can revert the tape to the original configuration \mathbf{w} (and then $\mathcal{D}(\bar{B}, V)$ simply issues the REVERT command). No matter the present configuration, we iterate through $I^0(V)$ in descending order. Letting the current index be $b \in I^0(V)$, recall this block is of form $(\mathbf{w}')^b = (z||*)$. First write $0^{50 \log(w/\varepsilon)}$ to the last indices (in reverse order to satisfy reversibility), such that we reach the configuration after compressing the block. Then enumerate over $h \in \{0, 1\}^{2t_0}$ using workspace $O(t_0 + \log(w/\varepsilon))$, until we find the hash with index z in \mathbf{BAD}_i , where \mathbf{BAD}_i and B_{t_0, \vec{h}_p} are defined as before (which we still have access to because we reset the tape in reverse order), and we determine membership by Proposition 3.11.

Once we find this h , write $(\mathbf{w}')^b = h$ (in the reverse order to satisfy reversibility) and proceed to the next highest index in $I^0(V)$. After this process has completed, it is clear from construction that \mathbf{w} has been reset to the original configuration, and that the tape never reaches a new intermediate configuration during this process.

Time and Space. In every step of the computation, we perform at most $\text{poly}(2^{t_0}nw/\varepsilon)$ work to determine if a hash function is good, find the index of a bad hash function, or find a good hash function. Moreover, as at every point we store at most a constant number of hash functions on the worktape, the space consumption follows. \square

It is easy to go from Theorem 3.10 to Theorem 3.1.

Proof of Theorem 3.1. Let $\overline{B} : [\text{poly}(n)] \times \{0, 1\}^{nc} \rightarrow [\text{poly}(n)]$ be the ROBP obtained from applying Lemma 3.2 to M with $n = n$, $w = n$, and $\varepsilon = 1/2n$. We then call Theorem 3.10 with $\overline{B} = \overline{B}$ and $\varepsilon = 1/2n^2$. Let $\vec{h} = (h_1, \dots, h_{c \log n})$ be the hash functions obtained from this call, which we have implicit access to via the current state of the catalytic tape \mathbf{w}' and V , and let $t = O(\log n)$ be the domain of the hash functions. Then enumerate over $x \in \{0, 1\}^t$ and for $i, j \in [w]$ let

$$\widetilde{M}_{i,j} = \Pr_{x \leftarrow U_t} [\overline{B}[i, \text{NIS}_{\vec{h}}(x)] = j].$$

By Lemma 3.2 and Theorem 3.10, we have the guarantee that

$$\|\widetilde{M} - M^n\|_1 \leq \frac{1}{2n} + n \cdot \frac{1}{2n^2} = 1/n. \quad \square$$

4 Catalytic Recursive Matrix Powering

We now transform Theorem 3.10 into a parameterized algorithm for matrix powering.

Theorem 4.1. *There is a catalytic machine that, given r_1, r_2 such that $r_1 r_2 = \log(n)$ and a stochastic matrix $M \in [0, 1]^{n \times n}$ where each entry is specified with $l = O(\log n)$ bits of precision, uses workspace $O(r_2 \cdot \log(n))$, catalytic space $O(r_1 \cdot \log(n))$, and time $n^{O(r_2)}$, and outputs \widetilde{M} such that $\|\widetilde{M} - M^n\| \leq 1/n$.*

Theorem 4.1 immediately implies Theorem 1.4 by setting $r_2 = \log^{\alpha-1}(n)$ and $r_1 = \log^{2-\alpha}(n)$ for $\alpha \in [1, 1.5]$, and using the standard transformation of estimating the acceptance probability of a **BPL** machine via stochastic matrix powering.

We first prove there exists an algorithm which computes a single intermediate power. We must be careful to ensure that the algorithm satisfies the requirements of (partial) catalytic computation. In particular, if the machine ever outputs an answer (rather than \perp), this must be the only possible answer for this input, over all possible catalytic tapes. Simultaneously, we must ensure that the vast majority of inputs never return \perp no matter the initial catalytic tape configuration.

We achieve this dual guarantee using an idea from Saks and Zhou [SZ99]. For a given input matrix M , we additionally take in a shift $s \in [0, \delta]$ for $\delta = 1/\text{poly}(w/\varepsilon)$. After computing an approximate power of M , we add s to each entry, and then truncate each entry to $O(\log w/\varepsilon)$ bits of precision. In fact, we first verify that our shifted approximate power is sufficiently far from the rounding threshold, and if not return \perp . By doing so, we algorithmically verify that we will never round to different thresholds over different \mathbf{w} . Unfortunately, for some pairs (M, s) it may be the case that we detect possible mis-rounding for some tapes \mathbf{w} , even if all possible approximations lie inside a single rounding interval. This can result in returning \perp on some tapes and a (consistent) value otherwise. However, we show that we can choose the magnitude of s such that with high probability over s this does not occur, and we always return a (consistent) value.

Theorem 4.2. *For every $n, w \in \mathbb{N}$ and $\varepsilon > 0$ where $w \geq n \geq \log w$ and $2^{-n} > \varepsilon > 0$, there is a reversible catalytic machine \mathcal{P} that uses workspace $O(\log w/\varepsilon)$, catalytic space $O(\log(n) \log(nw))$, and time $\text{poly}(nw/\varepsilon)$. The machine takes input $s \in \{0, 1\}^{O(\log(w/\varepsilon))}$ and a substochastic matrix $M \in [0, 1]^{w \times w}$, where each entry of M is specified with $l = O(\log w/\varepsilon)$ bits of precision. Moreover:*

- *For every (s, M) , there is substochastic \widetilde{M}_s (defined without reference to \mathbf{w}) with l bits of precision satisfying $\|\widetilde{M}_s - M^n\|_1 \leq \varepsilon$ such that for every \mathbf{w} ,*

$$\mathcal{P}^{\mathbf{w}}(s, M) \in \{\perp, \widetilde{M}_s\}.$$

- For every M , $\Pr_s[\exists \mathbf{w}, \mathcal{P}^{\mathbf{w}}(s, M) = \perp] \leq 1/w^2$.

Proof. Let $\overline{B} : [(w/\delta)^c] \times \{0, 1\}^{m=n \cdot O(\log(w/\delta))} \rightarrow [(w/\delta)^c]$ be the result of Lemma 3.2 applied to M with $n = n$ and $\varepsilon = \delta$ to be chosen later. We compose the output of this algorithm with Theorem 3.10, applied with $\overline{B} = \overline{B}$ and $w' = (w/\delta)^c$ and $n' = m = O(n^3)$ and $\varepsilon = \delta$ to be chosen later. Let $\vec{h} = (h_1, \dots, h_\ell)$ be the hash functions obtained from this call, which we have implicit access to via the current state of the catalytic tape \mathbf{w}' and V , and let $t = O(\log w/\delta)$ be the domain of the hash functions. Then enumerate over $x \in \{0, 1\}^t$ and for every $i, j \in [w]$ let

$$\widetilde{M}_{i,j} = \Pr_{x \leftarrow U_t} [\overline{B}[i, \text{NIS}_{\vec{h}}(x)] = j].$$

Next, define $\tau = (s \cdot 2^{-2k}) \cdot J$ where $J = 1^{w \times w}$, interpreting $s \in [2^k]$. We next check if any entry of $\widetilde{M} + \tau$ is within 2δ of a multiple of 2^{-l} , our rounding threshold. In this case, run $\mathcal{D}^{\mathbf{w}'}(\overline{B}, V)$ to reset the tape and return \perp . Otherwise, let

$$\widetilde{M}_s = \lfloor \widetilde{M} + \tau \rfloor_l$$

where $\lfloor \cdot \rfloor_l$ rounds each entry down to l bits of precision, and decreases the largest entry per row such that the final matrix is substochastic. Let this matrix be \widetilde{M}_s . Finally, run $\mathcal{D}^{\mathbf{w}'}(\overline{B}, V)$, and return \widetilde{M}_s .

Accuracy. By our choice of error in Theorem 3.10 and Lemma 3.2, we have that

$$\left\| \widetilde{M} - M^n \right\|_1 \leq 2w\delta$$

and moreover \widetilde{M} has each row sum at most 1. Furthermore, perturbing by τ and rounding down the largest entry causes an ℓ_1 error of at most $2w \cdot 2^{-k}$. Finally, rounding each entry down to a multiple of 2^{-l} causes a total error of at most $w \cdot 2^{-l}$, so

$$\left\| \widetilde{M}_s - M^n \right\| \leq 2w\delta + 2w \cdot 2^{-k} + w \cdot 2^{-l} \leq \varepsilon$$

Where the final inequality comes from choosing

$$l = O(\log(w/\varepsilon)), \quad k = 10 \cdot l, \quad \delta = 2^{-2k}.$$

Uniqueness. We claim that for every (s, M) , there is at most 1 possible non- \perp output over all choices of \mathbf{w} (which we denote \widetilde{M}_s in the theorem statement). Let $\widetilde{M}_{\mathbf{w}}, \widetilde{M}_{\mathbf{w}'}$ be the result of Theorem 3.10 on M initialized with catalytic tapes \mathbf{w}, \mathbf{w}' . By the accuracy guarantee of Theorem 3.10, for every i, j we have

$$\left| (\widetilde{M}_{\mathbf{w}} + \tau)_{i,j} - (\widetilde{M}_{\mathbf{w}'} + \tau)_{i,j} \right| \leq \left| (\widetilde{M}_{\mathbf{w}} + \tau)_{i,j} - (M + \tau)_{i,j} \right| + \left| (M + \tau)_{i,j} - (\widetilde{M}_{\mathbf{w}'} + \tau)_{i,j} \right| \leq 2\delta$$

Thus, if $(\widetilde{M}_{\mathbf{w}'} + \tau)_{i,j}$ is greater than 2δ from a multiple of 2^{-l} , we can be certain that no tape \mathbf{w}' will induce an estimate that falls on the other side of the threshold, and hence all non- \perp outputs will be rounded consistently.

Success Probability. Furthermore, we argue that for every M , with probability at least $1 - 1/w^2$ over s we return \widetilde{M}_s (not \perp) for every initial tape configuration. Fixing arbitrary s and $i, j \in [w]$, if $(M + \tau)_{i,j}$ is at least 3δ from every multiple of 2^{-l} , every \mathbf{w} will induce an estimate $(\widetilde{M} + \tau)_{i,j}$ that is at least 2δ from every multiple of 2^{-l} , and hence for every \mathbf{w} we will not produce \perp due to this entry. This occurs for every i, j simultaneously with probability at least

$$w^2 \cdot 2^l \cdot \frac{6\delta \cdot 2^k + 2}{2^k} \ll 1/w^2.$$

Time and Space. It is clear the algorithm runs in the claimed time and space bound, given Theorem 3.10.

Reversibility. As the only components of the algorithm that write to the catalytic tape are calls to Theorem 3.10, reversibility follows immediately from the equivalent result for that algorithm. \square

We can then prove the main result.

Theorem 4.1. *There is a catalytic machine that, given r_1, r_2 such that $r_1 r_2 = \log(n)$ and a stochastic matrix $M \in [0, 1]^{n \times n}$ where each entry is specified with $l = O(\log n)$ bits of precision, uses workspace $O(r_2 \cdot \log(n))$, catalytic space $O(r_1 \cdot \log(n))$, and time $n^{O(r_2)}$, and outputs \widetilde{M} such that $\|\widetilde{M} - M^n\| \leq 1/n$.*

Proof. Let $\vec{s} = (s_1, \dots, s_{r_2}) \in \{0, 1\}^{r_2 \cdot O(\log n)}$ be a vector of random shifts. Let $\widetilde{M}_0 = M$ and for $i \in [r_2]$ recursively define

$$\widetilde{M}_i = f\left(\widetilde{M}_{i-1}, s_i\right),$$

where f is the function defined by Theorem 4.2 with $\varepsilon = 1/n^3$ and $n = 2^{r_1}$. An easy inductive proof [PP23] establishes that, letting

$$\left\| \widetilde{M}_i - M^{2^{r_1 \cdot i}} \right\|_1 = \delta_i$$

we have $\delta_{i+1} \leq 1/n^3 + 2^{r_1} \cdot \delta_{i-1} \leq 2^{r_1+1} \cdot \delta_{i-1}$, and hence $\delta_{r_2} \leq 1/n$.

The final algorithm iterates over \vec{s} and computes \widetilde{M}_{r_2} by applying recursive composition of space-bounded machines Corollary 2.5 to the algorithm of Theorem 4.2 as defined above.⁷ The algorithm returns the first non- \perp output. The fact that the algorithm is catalytic follows from Corollary 2.5 and Theorem 4.2. Next, we claim there is some \vec{s} where the algorithm returns a value. Note that s_i is chosen obliviously to \widetilde{M}_{i-1} , and so with probability at least $1/n^2$ over s_i , on input $(\widetilde{M}_{i-1}, s_i)$ the algorithm returns \widetilde{M}_i (i.e. not \perp) when run with every possible catalytic tape. Thus, there is some \vec{s} where every level computes correctly.

Time and Space. Every application of Theorem 4.2 occurs with parameters $n = 2^{r_1}$ and $w = n$ and $\varepsilon = 1/n^3$, such that the algorithm uses workspace $O(r_1 + \log(n)) = O(\log n)$, catalytic space $O(r_1 \cdot (r_1 + \log(n))) = O(r_1 \cdot \log n)$, and time $\text{poly}(n)$, and moreover the shift s for each level is of length $O(\log n)$. Applying Corollary 2.5, we obtain that the composed algorithm uses workspace $O(r_2 \cdot \log(n) + |\vec{s}|) = O(r_2 \cdot \log n)$, catalytic space $O(r_1 \cdot \log n)$, and runs in time $n^{O(r_2)}$ as claimed. \square

⁷We can define the machine to take the entire shift vector and a pointer to the index it should use, such that we are recursively composing exactly the same machine, but we suppress this for clarity.

5 Acknowledgements

I thank Dean Doron, Roei Tell, and Ryan Williams for helpful discussions and comments on the manuscript. I thank Ian Mertz for helpful discussions about catalytic computing.

References

- [BCG20] Mark Braverman, Gil Cohen, and Sumegha Garg. Pseudorandom pseudo-distributions with near-optimal error for read-once branching programs. *SIAM J. Comput.*, 49(5), 2020.
- [BCK⁺14] Harry Buhrman, Richard Cleve, Michal Koucký, Bruno Loff, and Florian Speelman. Computing with a full memory: catalytic space. In *Symposium on Theory of Computing, STOC 2014*, pages 857–866. ACM, 2014.
- [BDS22] Sagar Bisoyi, Krishnamoorthy Dinesh, and Jayalal Sarma. On pure space vs catalytic space. *Theor. Comput. Sci.*, 921:112–126, 2022.
- [BKLS18] Harry Buhrman, Michal Koucký, Bruno Loff, and Florian Speelman. Catalytic space: Non-determinism and hierarchy. *Theory Comput. Syst.*, 62(1):116–135, 2018.
- [BRRY14] Mark Braverman, Anup Rao, Ran Raz, and Amir Yehudayoff. Pseudorandom generators for regular branching programs. *SIAM J. Comput.*, 43:973–986, 2014.
- [CDR⁺21] Gil Cohen, Dean Doron, Oren Renard, Ori Sberlo, and Amnon Ta-Shma. Error reduction for weighted prgs against read once branching programs. In *36th Computational Complexity Conference, CCC 2021, July 20-23, 2021*, volume 200 of *LIPICs*, pages 22:1–22:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [CDST23] Gil Cohen, Dean Doron, Ori Sberlo, and Amnon Ta-Shma. Approximating iterated multiplication of stochastic matrices in small space. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023*, pages 35–45. ACM, 2023.
- [CH22] Kuan Cheng and William M. Hoza. Hitting sets give two-sided derandomization of small space. *Theory of Computing*, 18(21):1–32, 2022.
- [CHHL18] Eshan Chattopadhyay, Pooya Hatami, Kaave Hosseini, and Shachar Lovett. Pseudorandom generators from polarizing random walks. In *33rd Computational Complexity Conference, CCC 2018*, volume 102 of *LIPICs*, pages 1:1–1:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [CHL⁺23] Lijie Chen, William Hoza, Xin Lyu, Avishay Tal, and Hongxun Wu. Weighted pseudorandom generators via inverse analysis of random walks and shortcutting. *Electron. Colloquium Comput. Complex.*, TR23-114, 2023.
- [CL20] Eshan Chattopadhyay and Jyun-Jie Liao. Optimal error pseudodistributions for read-once branching programs. In *35th Computational Complexity Conference, CCC 2020, July 28-31, 2020*, volume 169 of *LIPICs*, pages 25:1–25:27. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [CL23] Eshan Chattopadhyay and Jyun-Jie Liao. Recursive error reduction for regular branching programs. *CoRR*, abs/2309.04551, 2023.

- [CM20] James Cook and Ian Mertz. Catalytic approaches to the tree evaluation problem. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020*, pages 752–760. ACM, 2020.
- [DGJ⁺20] Samir Datta, Chetan Gupta, Rahul Jain, Vimal Raj Sharma, and Raghunath Tewari. Randomized and symmetric catalytic computation. In *Computer Science - Theory and Applications - 15th International Computer Science Symposium in Russia, CSR 2020*, pages 211–223. Springer, 2020.
- [DMR⁺21] Dean Doron, Raghu Meka, Omer Reingold, Avishay Tal, and Salil P. Vadhan. Pseudorandom generators for read-once monotone branching programs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2021*, volume 207 of *LIPICs*, pages 58:1–58:21, 2021.
- [DPT23] Dean Doron, Edward Pyne, and Roei Tell. Personal communication, 2023.
- [DT23] Dean Doron and Roei Tell. Derandomization with minimal memory footprint. In *38th Computational Complexity Conference, CCC 2023*, volume 264 of *LIPICs*, pages 11:1–11:15, 2023.
- [FK18] Michael A. Forbes and Zander Kelley. Pseudorandom generators for read-once branching programs, in any order. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018*, pages 946–955. IEEE Computer Society, 2018.
- [GJST19] Chetan Gupta, Rahul Jain, Vimal Raj Sharma, and Raghunath Tewari. Unambiguous catalytic computation. In *39th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2019*, volume 150 of *LIPICs*, pages 16:1–16:13, 2019.
- [Hoz21] William M. Hoza. Better pseudodistributions and derandomization for space-bounded computation. In *Proceedings of the 25th International Conference on Randomization and Computation (RANDOM)*, pages 28:1–28:23, 2021.
- [HZ18] William Hoza and David Zuckerman. Simple optimal hitting sets for small-success RL. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 59–64, 2018.
- [INW94] Russell Impagliazzo, Noam Nisan, and Avi Wigderson. Pseudorandomness for network algorithms. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994*, pages 356–364. ACM, 1994.
- [KvM02] Adam R. Klivans and Dieter van Melkebeek. Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses. *SIAM J. Comput.*, 31(5):1501–1526, 2002.
- [Mer23] Ian Mertz. Reusing space: Techniques and open problems. *Bulletin of EATCS*, 141(3), 2023.
- [Nis90] Noam Nisan. Pseudorandom generators for space-bounded computation. In Harriet Ortiz, editor, *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*, pages 204–212. ACM, 1990.

- [Nis93] Noam Nisan. On read-once vs. multiple access to randomness in logspace. *Theor. Comput. Sci.*, 107(1):135–144, 1993.
- [Nis94] Noam Nisan. $RL \leq SC$. *Comput. Complex.*, 4:1–11, 1994.
- [NZ96] Noam Nisan and David Zuckerman. Randomness is linear in space. *Journal of Computer and System Sciences*, 52(1):43–52, February 1996.
- [PP23] Aaron (Louie) Putterman and Edward Pyne. Near-optimal derandomization of medium-width branching programs. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023*, pages 23–34, 2023.
- [PRZ23] Edward Pyne, Ran Raz, and Wei Zhan. Certified hardness vs. randomness for log-space. *Electron. Colloquium Comput. Complex.*, TR23-040, 2023.
- [PV21] Edward Pyne and Salil Vadhan. Pseudodistributions That Beat All Pseudorandom Generators (Extended Abstract). In *36th Computational Complexity Conference (CCC 2021)*, pages 33:1–33:15, 2021.
- [Sav70] Walter J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4:177–192, 1970.
- [SZ99] Michael E. Saks and Shiyu Zhou. $BP_HSpace(S) \subseteq DSPACE(S^{3/2})$. *J. Comput. Syst. Sci.*, 58(2):376–403, 1999.

A Proofs of Lemmas

We first prove that a hash function drawn from a pairwise independent hash family is good for a function with high probability. To do this, we recall the hash mixing lemma:

Lemma A.1 ([Nis90]). *Let $A, A' \subseteq \{0, 1\}^t$ be arbitrary subsets of density $\rho = |A|/2^t$ and $\rho' = |A'|/2^t$. Then for every $\delta > 0$,*

$$\Pr_{h \leftarrow \mathcal{H}} \left[\left| \Pr_{x \leftarrow U_t} [x \in A, h(x) \in A'] - \rho\rho' \right| \geq \delta \right] \leq (1/\delta^2)/2^t.$$

Proof of Lemma 3.8. For every $i, k \in [w]$, let $A_{i,k} = \{x \in \{0, 1\}^t : f[i, x] = k\}$ and let $\rho_{i,k} = |A_{i,k}|/2^t$. Note that for every i, j ,

$$\Pr_{x, x' \leftarrow U_t} [f[f[i, x], x'] = j] = \sum_{k \in [w]} \rho_{i,k} \rho_{k,j}.$$

Thus, for every h such that for every i, k, j , $\Pr_{x \leftarrow U_t} [x \in A_{i,k}, h(x) \in A_{k,j}] - \rho_{i,k} \rho_{k,j} \leq \delta/w$, we have that h is δ -good for f . By Lemma A.1 this event occurs with probability $1 - (w/\delta)^2/2^t$ over $h \leftarrow \mathcal{H}$ for each pair of sets, and thus with probability $1 - w^3(w/\delta)^2/2^t = 1 - w^5(1/\delta)^2/2^t$ for every tuple (i, j, k) . \square

We recall there is a logspace algorithm which tests if a hash function is good, given oracle access to the function we wish to fool. We remark that there is work [Nis93, CH22, PRZ23] on testing if an entire PRG is good for a branching program, but we need a much weaker claim.

Lemma A.2 ([Nis94]). *There is a space $O(t + \log(w/\delta))$ algorithm that, given oracle access to $f : [w] \times \{0, 1\}^t \rightarrow [w]$ and $h \in \mathcal{H}_t$ and $\delta > 0$, tests if h is δ -good for f .*

Proof. The algorithm enumerates over $i, j \in [w]$. For every i, j , the algorithm computes $p_{i,j} = \mathbb{E}_{x,x' \leftarrow U_t}[f[f[i,x],x']]$ (i.e. the correct probability) by enumeration over x, x' in space $O(t + \log w)$. Then it computes $\tilde{p}_{i,j} = \mathbb{E}_{x \leftarrow U_t}[f[f[i,x],h(x)]]$ and rejects if the estimate is greater than δ from the true value. Correctness and total space consumption are immediate. \square

We can then prove that we can test if a hash function is good, given \overline{B} and pointers to preceding hash functions.

Proposition 3.11. *There is a space $O(t + \log(w/\delta))$ algorithm that, given $n, w, t \in \mathbb{N}$ with $w \geq n$ and $\tilde{h} \in \{0, 1\}^{2t}$ and $\overline{B} : [w] \times \{0, 1\}^n \rightarrow [w]$ and (read only) \mathbf{w} and pointers p_1, \dots, p_r such that $\mathbf{w}_{[p_i, \dots, p_i+2t]} = h_i$ represents a hash function on t bits, returns if \tilde{h} is δ -good for $B_{t,(h_1, \dots, h_r)}$.*

Proof. By Lemma A.2, it suffices to show that given $i \in [w]$ and $x \in \{0, 1\}^t$, we can compute $B_{t,(h_1, \dots, h_r)}[i, x]$. To do this, the algorithm maintains $v \in [w]$ as its current position in the branching program (initialized to $v = i$) and $i \in [n]$ to track the current layer. To determine the next position, it suffices to determine the i th block of the output of $\text{NIS}_{(h_1, \dots, h_r)}(x)$. It is well known that this can be computed in space $O(t + \log nw)$ given read-only access to the set of hash functions (by walking down the binary expansion of i , denoted $\langle i \rangle$, and applying h_j if $\langle i \rangle_j = 1$), which we have via the pointers. \square

Finally, we provide a translation of our quantization statement. We first recall a strict specialization of the statement of [PP23]:

Lemma A.3. *There exists a **canonicalizer** algorithm \mathbf{C}_t that, given $n, w \in \mathbb{N}$ with $w \geq n$, takes in $\varepsilon > 0$ and a sub-stochastic matrix $M \in \mathbb{R}^{w \times w}$ with each entry represented by at most $O(\log w/\varepsilon)$ bits, runs in space $O(\log w/\varepsilon)$, and returns a branching program \overline{B} of length n and width $w + 1$ with alphabet $\{0, 1\}^t$ for $t = O(\log(w/\varepsilon))$. Moreover, letting $\widetilde{M} \in [0, 1]^{w \times w}$ be the matrix where for $i, j \in [w]$ we have*

$$\widetilde{M}_{i,j} = \Pr_{x \leftarrow U_{(\{0,1\}^t)^n}} [\overline{B}[i, x] = j]$$

then

$$\left\| \widetilde{M} - M^n \right\|_1 \leq \varepsilon.$$

We reduce the alphabet (and slightly increase the length) as follows. We transform \overline{B} into a branching program of width $(w + 1) \cdot 2^t = \text{poly}(w/\varepsilon)$ and length $n \cdot t = n \cdot O(\log w/\varepsilon)$, where each set of t layers reads t bits, interprets these bits as $\sigma \in \{0, 1\}^t$, and takes the transition labeled with σ in \overline{B} .