



Theory of Unconditional Pseudorandom Generators

Pooya Hatami
CSE Department
The Ohio State University
pooyahat@gmail.com

William M. Hoza
Simons Institute
University of California, Berkeley
williamhoza@berkeley.edu

Abstract

This is a survey of unconditional *pseudorandom generators* (PRGs). A PRG uses a short, truly random seed to generate a long, “pseudorandom” sequence of bits. To be more specific, for each restricted model of computation (e.g., bounded-depth circuits or read-once branching programs), we would like to design a PRG that “fools” the model, meaning that every function computable in the model behaves approximately the same when we plug in pseudorandom bits from the PRG as it does when we plug in truly random bits. In this survey, we discuss four major paradigms for designing PRGs:

- We present several PRGs based on k -wise uniform generators, small-bias generators, and simple combinations thereof, including proofs of Viola’s theorem on fooling low-degree polynomials (Comput. Complexity 2009) and Braverman’s theorem on fooling \mathbf{AC}^0 circuits (J. ACM 2010).
- We present several PRGs based on “recycling” random bits to take advantage of communication bottlenecks, such as the Impagliazzo-Nisan-Wigderson generator (STOC 1994).
- We present connections between PRGs and computational hardness, including the Nisan-Wigderson framework for converting a hard Boolean function into a PRG (J. Comput. Syst. Sci. 1994).
- We present PRG frameworks based on random restrictions, including the “polarizing random walks” framework (Chattopadhyay, Hatami, Hosseini, and Lovett, Theory Comput. 2019).

We explain how to use these paradigms to construct PRGs that work *unconditionally*, with no unproven mathematical assumptions. The PRG constructions use ingredients such as finite field arithmetic, expander graphs, and randomness extractors. The analyses use techniques such as Fourier analysis, sandwiching approximators, and simplification-under-restrictions lemmas.

Acknowledgments We thank Yevgeniy Dodis, Avishay Tal, and David Zuckerman for helpful comments on a draft of this work.

Contents

- 1 Introduction** **4**
- 1.1 Whom shall we fool? Three PRG paradigms 5
 - 1.1.1 PRGs for everyday non-adversarial applications 5
 - 1.1.2 PRGs for all efficient observers 6
 - 1.1.3 PRGs for restricted models of computation 7
- 1.2 Overview of this text 8
- 1.3 The generic probabilistic existence proof 9
- 1.4 Explicitness 10
 - 1.4.1 Families of PRGs 10
 - 1.4.2 The default conjecture: Explicit PRGs exist 11
- 1.5 Beyond PRGs: Hitting set generators and more 12

- 2 Limited Independence and Small-Bias Generators** **14**
- 2.1 Limited independence 14
 - 2.1.1 Pairwise uniform bits 14
 - 2.1.2 k -wise uniform bits 15
 - 2.1.3 Perfectly fooling shallow decision trees 16
 - 2.1.4 Connection with coding theory: Dual codes 17
- 2.2 Small-bias distributions 18
 - 2.2.1 Fooling parities of variables 18
 - 2.2.2 A better seed length for parities of few variables 19
 - 2.2.3 Connection with coding theory: Nearly balanced codes 20
- 2.3 Analysis technique: Fourier L_1 bounds 21
 - 2.3.1 Basic Fourier analysis 21
 - 2.3.2 Almost k -wise uniform bits 21
 - 2.3.3 Fooling bounded-size decision trees 22
 - 2.3.4 Fooling width-2 branching programs 24
- 2.4 Viola’s generator for low-degree \mathbb{F}_2 -polynomials 25
 - 2.4.1 Directional derivatives 26
 - 2.4.2 The XOR of two independent copies of an arbitrary Boolean function 27
 - 2.4.3 The reduction from fooling \mathcal{F} to fooling $\partial\mathcal{F}$ 27
 - 2.4.4 Inductive analysis of low-degree polynomials 29
 - 2.4.5 Application: Width-2 branching programs that read several bits at a time 29
- 2.5 Analysis technique: Sandwiching approximators 30
 - 2.5.1 The sandwiching lemma 30
 - 2.5.2 Using k -wise uniform generators to fool size- m decision trees 31
 - 2.5.3 Small-bias distributions fool read-once \mathbf{AC}^0 31

2.5.4	The sandwiching lemma and the triangle inequality are always enough	34
2.6	Braverman’s theorem: Limited independence fools \mathbf{AC}^0	37
2.6.1	LMN polynomials	38
2.6.2	Operations on functions with low-degree sandwiching polynomials	38
2.6.3	Low-degree sandwichers for \mathbf{AC}^0 circuits	39
2.6.4	Improved parameters via probabilistic polynomials	41
3	Recycling Random Bits	45
3.1	PRGs for two-party communication protocols	45
3.1.1	Expander graphs from a PRG perspective	45
3.1.2	Combinatorial rectangles and the Expander Mixing Lemma	47
3.2	The INW generator for standard-order ROBPs	48
3.2.1	Concatenating two independent pseudorandom strings	49
3.2.2	Recycling seeds using a PRG for two-dimensional rectangles	50
3.3	The BRRY generator for standard-order regular ROBPs	51
3.3.1	Improved analysis of the INW generator for low-weight programs	51
3.3.2	Regular programs have low weight	55
3.4	The Nisan-Zuckerman generator for short, wide ROBPs	56
3.4.1	Randomness extractors	56
3.4.2	Using extractors to fool standard-order ROBPs	58
4	PRGs and Hardness	61
4.1	PRGs as high-quality lower bounds	61
4.1.1	PRGs imply hard Boolean functions	61
4.1.2	The lack-of-lower-bounds barrier	63
4.2	The Nisan-Wigderson framework	64
4.2.1	Constructing a PRG from a hard function	64
4.2.2	Analysis: Unpredictability	64
4.2.3	A family of nearly disjoint sets	66
4.2.4	Unconditional applications	67
4.3	Hardness-based PRGs beyond Nisan-Wigderson	68
5	Random Restrictions	70
5.1	PRGs from polarizing random walks	72
5.1.1	Simplification under truly random restrictions	72
5.1.2	Fractional PRGs	72
5.1.3	From fractional PRGs to PRGs	74
5.1.4	A better reduction for the low-error regime	78
5.2	Analysis technique: Fourier growth bounds	80
5.2.1	The noise operator and simplification on average	80
5.2.2	Fourier growth bounds for regular ROBPs	81
5.2.3	Using Fourier growth bounds to obtain PRGs	85
5.3	Fooling \mathbf{AC}^0 via the Ajtai-Wigderson framework	88
5.3.1	Simplification under partially pseudorandom restrictions	88
5.3.2	Restrictions that preserve expectation	90
5.3.3	Iterated restrictions	91
5.4	The Forbes-Kelley generator for ROBPs	92
5.4.1	Pseudorandomness plus noise	92

5.4.2	A Fourier decomposition lemma for ROBPs	93
5.4.3	Pseudorandom restrictions that preserve the expectation of ROBPs	95
5.4.4	A better generator for the small-width setting	97
5.5	PRGs for read-once CNFs via early termination	97
5.5.1	Simplification of read-once CNFs under fully-pseudorandom restrictions	98
5.5.2	Iterated restrictions with early termination	100
5.5.3	Discussion: Two types of simplification	101
5.6	Fooling general branching programs via the IMZ framework	102
5.6.1	Shrinkage of branching programs under fully-pseudorandom restrictions	103
5.6.2	PRGs from fully-derandomized shrinkage lemmas	105
6	Table of PRGs	108
6.1	Circuit models	108
6.2	Branching program models	108
6.3	Algebraic models	110
6.4	Models based on locality	110

Chapter 1

Introduction

To make random choices, it would be useful to have an unlimited supply of “truly random” bits: unbiased and independent coin flips. What can we do if we only have a few truly random bits? A *pseudorandom generator* (PRG) uses a small amount of true randomness, called the “seed,” to generate a long sequence that appears to be completely random (even though it isn’t). PRGs are ubiquitous in computing theory and practice. The basic motivation is that we think of randomness as a scarce computational resource, akin to time or space, so whenever we get our hands on some random bits, we want to stretch them as far as possible.

To model PRGs mathematically, we consider some “observer,” modeled as a function f . Let U_n denote the uniform distribution over $\{0, 1\}^n$. We would like to “fool” f in the following sense.

Definition 1.0.1 (Fooling). *Suppose $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is a function, X is a probability distribution over $\{0, 1\}^n$, and $\varepsilon > 0$. We say that X fools f with error ε , or ε -fools f , if*

$$|\Pr[f(X) = 1] - \Pr[f(U_n) = 1]| \leq \varepsilon.$$

More generally, we can consider a real-valued function $f: \{0, 1\}^n \rightarrow \mathbb{R}$. In this case, we say that X fools f with error ε if

$$|\mathbb{E}[f(X)] - \mathbb{E}[f(U_n)]| \leq \varepsilon.$$

If $\varepsilon = 0$, we say that X perfectly fools f .

Remark 1.0.2. *As a shorthand, we often identify the function f with the random variable $f(U_n)$. For example, instead of $\mathbb{E}[f(U_n)]$, we simply write $\mathbb{E}[f]$.*

Definition 1.0.1 says that although X might not be uniform, X and U_n are nevertheless *indistinguishable*, at least from f ’s perspective. Conversely, if X does *not* ε -fool f , we refer to f as a “distinguisher” for X . A PRG’s job is to use a few truly random bits to sample a distribution that fools f .

Definition 1.0.3 (PRGs). *Suppose $f: \{0, 1\}^n \rightarrow \mathbb{R}$ and $G: \{0, 1\}^s \rightarrow \{0, 1\}^n$ are functions and $\varepsilon > 0$. We say that G is an ε -PRG for f if $G(U_s)$ fools f with error ε . In this case, we also say that G fools f with error ε . (See [Figure 1.1](#).)*

The parameter s is called the *seed length* of the PRG; we would like s to be as small as possible. Throughout this text, the parameter “ n ” will always denote the number of pseudorandom bits we are generating.

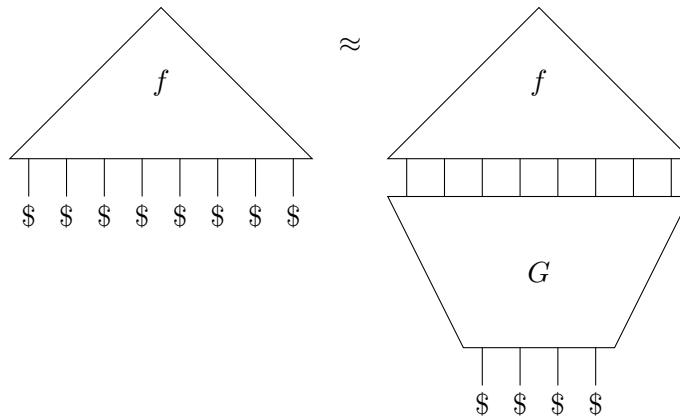


Figure 1.1: A PRG (G) uses a few truly random bits (depicted here using \$ symbols) to sample a pseudorandom string that is indistinguishable from a truly random string, from the perspective of the observer (f).

1.1 Whom shall we fool? Three PRG paradigms

An unavoidable fact of life is that for any nontrivial PRG, there exists a function that is not fooled by the PRG.

Claim 1.1.1 (Impossibility of fooling all functions). *Let $G: \{0, 1\}^s \rightarrow \{0, 1\}^n$ where $s < n$. There exists some $f: \{0, 1\}^n \rightarrow \{0, 1\}$ such that G does not 0.49-fool f .*

Proof. Let f be the indicator function for the image of G . Then $\mathbb{E}[f(G(U_s))] = 1$, whereas $\mathbb{E}[f] \leq 1/2$ because $s < n$. \square

In light of [Claim 1.1.1](#), the best we can hope for is generating bits that fool some large sets of observers but not all of them. After all, as Avi Wigderson says, *randomness is in the eye of the beholder*. If \mathcal{F} is a class of functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$, we say that G is an ϵ -PRG for \mathcal{F} if G ϵ -fools every $f \in \mathcal{F}$.

Which observers shall we fool? The study of PRGs can be crudely divided into three paradigms based on three possible answers:

1. Everyday non-adversarial applications.
2. All efficient observers.
3. Restricted models of computation.

We discuss these three paradigms in [Sections 1.1.1](#) to [1.1.3](#).

1.1.1 PRGs for everyday non-adversarial applications

In practice, when programmers want randomness, they invoke some type of `random()` method provided by the computing environment. Under the hood, these `random()` methods typically involve several components, each of which might be quite sophisticated. When practitioners speak of “pseudorandom number generators” or “random number generators,” they are usually referring to the entire randomness system as a whole, including whatever techniques are used to produce an initial seed. For example, the system might derive a seed from the current time of day, even though

such a seed is rather predictable. As another example, the system might use hardware random number generators based on thermal noise measurements.

In this text, we sidestep the important issue of producing a seed, along with many other issues that are important in practice. We focus on the challenge of stretching a truly random seed out to a long pseudorandom string. In our terminology, this is the job of a PRG (see [Definition 1.0.3](#)). A PRG is thus one of multiple components of a practical randomness system. For example, Java’s `Math.random()` method currently uses a type of PRG called a *linear congruential generator*. For such a PRG, the seed is a random number $X_0 \in \{0, 1, \dots, M - 1\}$, and the output sequence is (X_1, X_2, X_3, \dots) , where

$$X_{i+1} = a \cdot X_i + b \bmod M$$

for some parameters M, a, b . Meanwhile, Python’s `random.random()` method uses an algorithm called the “Mersenne twister” [\[MN98\]](#), and major web browsers currently use a PRG in the “xorshift+ family” [\[Vig17\]](#) to implement Javascript’s `Math.random()` function.

Why these PRGs are unsatisfactory

Practitioners use these randomness systems for both casual applications (e.g., video games) and serious applications (e.g., randomized algorithms). However, for a generic randomized algorithm, there is *no firm mathematical guarantee* that the outputs will be reliable when the algorithm is executed using one of these practical randomness systems. The methods that practitioners typically use to run randomized algorithms must be considered *heuristics*.

To be clear, a lot of work goes into designing high-quality practical randomness systems. Designers strive to ensure that these systems can be safely used in any application that “comes up naturally” in practice. The system is only deemed acceptable for everyday use when it passes a great number of creative statistical tests, such as those in the TestU01 family [\[LS07\]](#).

These statistical tests are valuable, but there is a wide gap between the statistical tests and a typical randomized algorithm. The designers behind practical systems such as Java’s `Math.random()` method wisely do not claim that they work in *adversarial* scenarios, so these systems are considered unsuitable for cryptography. This is true even if we focus solely on the PRG component of these systems. Furthermore, sometimes programs “accidentally” distinguish pseudorandom numbers from truly random numbers. There are quite a few documented cases in which PRGs have been shown to cause inaccurate scientific simulations [\[KW84; PR85; FMF85; MBH86; FLW92; Gra93; Cod94; CLK11\]](#)! One must imagine that other cases have gone unnoticed.

To a theoretician, this state of affairs is deeply unsatisfactory. Yes, modern practical PRGs seem to almost always work well in practice, but we don’t have a mathematically rigorous explanation for *why* these systems work. It’s not even clear what precisely the goal is. (Mathematically, how can we make a distinction between “adversarially-designed” programs and “naturally-occurring” programs?) By theoreticians’ standards, the success of practical PRGs is largely a mystery.

1.1.2 PRGs for all efficient observers

One of the great ideas in the theory of computing is to try to design a PRG that fools *all computationally efficient* observers. Given such a PRG and a truly random seed, we would be able to execute any randomized algorithm that is actually worth executing. (After all, there’s no point running a program if one won’t even survive long enough to see the output!) Such a PRG could also be used in cryptographic settings, because we can safely assume that eavesdroppers and hackers

only have so much computational power.¹

For example, given a random seed $X_0 \in \{1, 2, \dots, M-1\}$, the Blum-Blum-Shub (BBS) generator [BBS86] outputs the sequence $(X_1 \bmod 2, X_2 \bmod 2, X_3 \bmod 2, \dots)$ where

$$X_{i+1} = X_i^2 \bmod M.$$

This PRG is reminiscent of linear congruential generators, but the similarity is only superficial. For a suitably chosen modulus M , it is believed that the BBS generator fools all polynomial-time algorithms.

Why these PRGs are also (currently) unsatisfactory

Fooling all efficient observers is a well-defined and well-motivated *goal*. Unfortunately, nobody knows how to prove that some efficiently-computable PRG actually *has* this marvelous property.

To be clear, there is a substantial body of “evidence” indicating that such PRGs exist. For example, Blum, Blum, and Shub showed that their generator fools all polynomial-time observers, under the plausible-but-unproven assumption that there is no good algorithm for the “quadratic residuosity problem” [BBS86]. There are many other examples of PRGs that fool all polynomial-time observers under reasonable cryptographic or complexity-theoretic assumptions [Yao82; BM84; NW94; IW97; HILL99; KM02; Uma03; DMOZ20]. For practical cryptography, software developers tend to use PRGs that are not even supported by rigorous *conditional* proofs of correctness, but rather are supported by heuristic and intuitive arguments.

There is a genuine possibility that these PRGs are not secure. In one infamous incident, the U.S. National Institute of Standards and Technology (NIST) recommended using a PRG called “Dual_EC_DRBG.” The PRG was designed by the U.S. National Security Agency (NSA), and allegedly, they *intentionally designed it to be insecure* for surveillance purposes [Per13].

Once again, to a theoretician, this state of affairs is not satisfactory. There is genuine room for doubt about *whether* known PRGs work, and perhaps more importantly, even if they do work, we don’t have a good explanation for *why* they work. Conditional proofs can be considered partial explanations at best. The problem of designing PRGs that unconditionally fool all efficient observers is very challenging, with connections to deep topics such as the famous **P** vs. **NP** problem. (See Section 4.1.)

1.1.3 PRGs for restricted models of computation

The main topic of this text is a third paradigm for studying PRGs. In this third paradigm, we identify an interesting and well-defined *restricted model of computation*. Then we design PRGs that fool the chosen model of computation (unconditionally – with no unproven assumptions) and try to optimize the seed length of the PRG.

A toy example might clarify the idea. Let us design a PRG $G: \{0, 1\}^2 \rightarrow \{0, 1\}^3$ that fools every observer f that only looks at two of the three output bits. This problem is not completely trivial, because we don’t know which two bits f will observe. Nevertheless, the problem can be solved by defining

$$G(u_1, u_2) = (u_1, u_2, u_1 \oplus u_2),$$

¹There is a subtle distinction here. In the context of randomized algorithms, it’s okay if the PRG itself uses a little more time than the algorithms that we are trying to fool. On the other hand, in the context of cryptography, we want an efficiently-computable PRG that fools all efficient adversaries, including those that use polynomially more time than the PRG uses.

where \oplus denotes the XOR operation. When u_1 and u_2 are chosen uniformly at random, the three output bits are correlated, but any two of the bits are independent and uniform random.

Unconditional PRGs can be constructed for much richer and more interesting restricted models of computation. We are especially interested in fooling models of computation that have a “complexity theory” flavor, i.e., we want the output of the PRG to appear random to any observer that is “sufficiently efficient” in some sense. Arguably, the two most important models in this field are *constant-depth circuits* (\mathbf{AC}^0 , see [Definition 2.5.4](#)) and *read-once branching programs* (ROBPs, see [Definition 3.2.1](#)).

The value of these PRGs

Could PRGs for restricted models ever be directly used in practical applications? Potentially. PRGs for restricted models can be used to simulate some randomized algorithms. For example, suppose a randomized decision algorithm A uses S bits of space. For any fixed input a of A , we can consider the output of A as a function of its random bits, say $f(x) = A(a, x)$. It turns out that this function f can be computed by an RBP of width $2^{O(S)}$. Therefore, PRGs that fool ROBPs can be used to simulate A without significantly distorting its behavior.

Admittedly, it’s a bit unrealistic to imagine the PRGs studied in the theoretical literature being implemented on actual computers, because it is hard to compete with the practical PRGs discussed in [Section 1.1.1](#). Instead, the study of PRGs for restricted models has a much grander and broader purpose: these PRGs help to uncover the mysteries of the theory of computing, and hence are invaluable from a scientific perspective.

Unconditional PRGs have many applications within theoretical computer science, and we will not attempt to survey them here, but we will elaborate on one of the more important applications. A major open problem asks whether randomized decision algorithms have any *intrinsic* advantage over deterministic decision algorithms in terms of space complexity. PRGs for ROBPs have the potential to resolve this question. As discussed above, a PRG for ROBPs can be used to simulate randomized space-bounded decision algorithms using just a little bit of randomness. By exhaustively trying all seeds of the PRG and taking the majority outcome, we can actually get a completely deterministic simulation. An optimal PRG for ROBPs would imply that randomized space- S algorithms can be simulated deterministically in space $O(S)$, and hence randomness only confers a constant-factor advantage.

So far, optimal constructions of PRGs for ROBPs are not known, but we do have “pretty good” constructions (e.g., see [Section 3.2](#)). There are many partial derandomization results known for space-bounded computation, building on the theory of PRGs for ROBPs (in nontrivial ways). For example, it has been shown that randomized space- S algorithms can be simulated deterministically in space slightly less than $S^{3/2}$ [[SZ99](#); [Hoz21](#)]. There is no particular “barrier” known preventing us from designing optimal PRGs for ROBPs. This exciting problem is a central open problem in the unconditional theory of PRGs.

Apart from any application, we hope to convince the reader that PRGs for restricted models are interesting in their own right.

1.2 Overview of this text

In this work, we survey some of the most important frameworks and techniques for constructing unconditional PRGs for restricted models of computation. We focus on four major PRG paradigms:

- In [Chapter 2](#), we present k -wise uniform generators, small-bias generators, and simple combinations thereof.
- In [Chapter 3](#), we present PRGs that “recycle” randomness to take advantage of communication bottlenecks, such as the Impagliazzo-Nisan-Wigderson generator [[INW94](#)].
- In [Chapter 4](#), we present connections between PRGs and computational hardness, including the Nisan-Wigderson framework for converting a hard Boolean function into a PRG [[NW94](#)].
- In [Chapter 5](#), we present methods for constructing PRGs based on (pseudo)random restrictions, including the relatively recent “polarizing random walks” framework [[CHHL19](#)].

Along the way, as needed, we introduce the computational models that we fool (decision trees, circuits, branching programs, etc.) and techniques for analyzing PRGs (Fourier analysis, sandwiching approximators, simplification-under-restriction lemmas, etc.)

The literature on unconditional PRGs is vast, and this survey is far from exhaustive. (For example, we do not discuss the important line of work on fooling *linear threshold functions* [[RS10](#); [DGJSV10](#); [MZ13](#); [GKM18](#)].) Instead, we hope that this work serves as a suitable *introduction* to the field of unconditional PRGs, preparing the reader to study new and old papers on PRGs and make their own contributions.

The results that we cover include both classic and recent works. Besides covering the most important *principles* of PRG design and analysis, we also made sure to include expositions of many of the most important *examples* of unconditional PRGs, such as Viola’s PRG for low-degree polynomials [[Vio09](#)], Braverman’s theorem that limited independence fools \mathbf{AC}^0 [[Bra10](#)], and Forbes and Kelley’s relatively recent PRG for arbitrary-order ROBPs [[FK18](#)].

This text is primarily expository. However, we couldn’t help but include a few novel theorems and proofs. For example, we present a new proof of Braverman’s theorem ([Section 2.6](#)), and we present a new improvement to the polarizing random walks framework in the low-error regime ([Section 5.1.4](#)). We also highlight plenty of important open problems regarding PRGs for restricted models of computation.

Many wonderful prior expository works [[Mil01](#); [LW06](#); [AB09](#); [Gol10](#); [Vad12](#); [O’D14](#)] and lecture notes [[Zuc01](#); [Tre05](#); [Ta-15](#); [Ta-16](#); [Vio17](#); [Cha18](#); [Ta-18](#); [Cha19](#); [Ta-19](#); [Tal21](#); [Cha22](#)] include some coverage of unconditional PRGs. However, none of them have quite the same focus as our work, so we feel that our work fills a gap.

In the rest of this chapter, we discuss some additional basic issues related to the concept of a PRG, paving the way for the PRG constructions in subsequent chapters.

1.3 The generic probabilistic existence proof

For many classes \mathcal{F} , including classes defined by standard computational models (such as decision trees, circuits, branching programs, etc.), there is a totally generic argument showing that there exist PRGs that fool \mathcal{F} with a small seed length.

Proposition 1.3.1 (Nonexplicit PRGs). *Let \mathcal{F} be a class of functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$. For every $\varepsilon > 0$, there exists an ε -PRG for \mathcal{F} with seed length $\log \log |\mathcal{F}| + 2 \log(1/\varepsilon) + O(1)$.*

Proof. Pick a function $G: \{0, 1\}^s \rightarrow \{0, 1\}^n$ uniformly at random. Consider any arbitrary $f \in \mathcal{F}$. For each seed y , the value $f(G(y))$ is a random bit satisfying

$$\mathbb{E}_G[f(G(y))] = \mathbb{E}_{U_n}[f(U_n)].$$

Furthermore, as y ranges over all 2^s possible seeds, these random variables $f(G(y))$ are independent. Therefore, by Hoeffding’s inequality,

$$\Pr_G \left[\left| \mathbb{E}[f] - 2^{-s} \sum_{y \in \{0,1\}^s} f(G(y)) \right| > \varepsilon \right] \leq 2e^{-2\varepsilon^2 2^s}.$$

By the union bound, the probability that G fails to ε -fool \mathcal{F} is bounded by $2|\mathcal{F}|e^{-2\varepsilon^2 2^s}$. For $s = \log \log |\mathcal{F}| + 2 \log(1/\varepsilon) + O(1)$, this probability is less than 1, i.e., there exists a G that does ε -fool \mathcal{F} . \square

In a typical case – e.g., if \mathcal{F} is the set of all circuits of size at most n – each function $f \in \mathcal{F}$ can be described using $\text{poly}(n)$ bits, i.e., $|\mathcal{F}| \leq 2^{\text{poly}(n)}$. In this case, the PRG guaranteed by [Proposition 1.3.1](#) has seed length $O(\log(n/\varepsilon))$.

1.4 Explicitness

[Proposition 1.3.1](#) has a major weakness: it does not guarantee that the PRG is *efficiently computable*. The proof of [Proposition 1.3.1](#) is in some sense “nonconstructive.” Ideally, we want an *algorithm* for sampling from a pseudorandom distribution, and we want the algorithm to be reasonably efficient with respect to randomness *and* more conventional complexity measures simultaneously.

Definition 1.4.1 (Explicitness). *A PRG $G: \{0, 1\}^s \rightarrow \{0, 1\}^n$ is explicit if it can be computed in time $\text{poly}(n)$.*

One could consider alternative standards of explicitness. We could require that each individual output bit can be computed in time $\text{polylog } n$, or that the PRG runs in space $O(\log n)$, or that each bit can be computed in \mathbf{AC}^0 , or any number of other conditions. The truth is, there is no “one true definition” of explicitness. The appropriate definition depends on what one hopes to gain from the PRG.

For example, one might plan to derandomize an algorithm by *exhaustively trying all possible seeds* of a PRG. In this case, since we are inevitably going to spend more than 2^s steps on this brute-force process, it might make sense to relax our standard of “explicitness” and allow the PRG’s time complexity to be as large as $2^{O(s)} \cdot \text{poly}(n)$ rather than $\text{poly}(n)$. As another example, it turns out that PRGs can be used to prove that certain models of computation cannot solve the “Minimum Circuit Size Problem” [KC00]. For this application, the “correct” definition of explicitness is that for each fixed seed $u \in \{0, 1\}^s$, there is a Boolean circuit C_u of size $o(n/\log n)$ such that for every $i \in [n]$, we have $C_u(i) = G(u)_i$.

In this text, we will stick with [Definition 1.4.1](#) for concreteness, but when we present PRG constructions, we will generally not bother carefully verifying the runtime bound. Instead, we will focus on making the construction clear to the reader.

1.4.1 Families of PRGs

[Definition 1.4.1](#) refers to the time complexity of a PRG. To meaningfully speak of time complexity, we technically ought to be considering a whole *family* of PRGs. The convention in this line of work is to keep the family implicit. For example, a theorem might say something like the following.

For all $n, m \in \mathbb{N}$ and all $\varepsilon > 0$, there exists an explicit ε -PRG for size- m decision trees on n input bits with seed length $O(\log(m/\varepsilon) + \log \log n)$.

(See [Section 2.3.3](#).) Translating into more precise language, the same theorem can be restated as follows.

There exists a randomized algorithm \mathcal{G} satisfying the following.

1. Given input parameters n, m, ε , the algorithm \mathcal{G} outputs a string $\mathcal{G}(n, m, \varepsilon) \in \{0, 1\}^n$.
2. For all n, m, ε , the output distribution $\mathcal{G}(n, m, \varepsilon)$ fools size- m decision trees with error ε .
3. $\mathcal{G}(n, m, \varepsilon)$ uses at most $O(\log(m/\varepsilon) + \log \log n)$ random bits and runs in time $\text{poly}(n)$.

There is something potentially troubling about this “translation” process. The quantifiers got flipped! In the informal theorem statement, we say “for all n, m, ε , there exists an explicit PRG,” but strictly speaking, we mean that there exists a single algorithm \mathcal{G} that works for all n, m, ε simultaneously! Is this “flipped quantifiers” convention wise?

Let us make an analogy with big- O notation. Recall, e.g., the famous planar separator theorem:

For all $n \in \mathbb{N}$, for every n -vertex planar graph, there exists a set of $O(\sqrt{n})$ vertices such that removing those vertices splits the graph into connected components with at most $2n/3$ vertices each.

If we wanted to be more rigorous, we ought to flip the quantifiers and write something like the following:

There exists a function $f: \mathbb{N} \rightarrow \mathbb{N}$ such that $f \in O(\sqrt{n})$ and for all $n \in \mathbb{N}$, for every n -vertex planar graph, there exists a set of $f(n)$ vertices such that removing those vertices splits the graph into connected components with at most $2n/3$ vertices each.

We don’t bother with such careful language because it obscures more than it clarifies. The important thing is that the expression “ $O(\sqrt{n})$ ” tells how the number of removed vertices scales with the universally quantified parameter n . Analogously, when we say “there exists an explicit PRG,” the word “explicit” tells how the *computational complexity* of the PRG scales with the parameters.

1.4.2 The default conjecture: Explicit PRGs exist

For each “reasonable” class \mathcal{F} , the standard conjecture is that there exists an *explicit* PRG with essentially the same seed length as the generic nonexplicit bound ([Proposition 1.3.1](#)). Oftentimes, this conjecture can be supported with evidence in the form of conditional constructions. For example, consider the class \mathcal{F} of all CNF formulas of size at most n . The nonexplicit PRG has seed length $O(\log(n/\varepsilon))$. Under plausible complexity-theoretic assumptions, there is indeed an explicit PRG for *all* size- n Boolean circuits (whether CNF formulas or not) with seed length $O(\log(n/\varepsilon))$ [[IW97](#)].

Even without a compelling conditional construction, the “default” conjecture would be that a probabilistic existence proof can be matched by an explicit construction. The main challenge is to find the explicit construction. Typically, such a PRG would be optimal, i.e., one can unconditionally prove a seed length lower bound matching the nonexplicit bound to within a constant factor.² For example, every PRG for size- n CNF formulas (explicit or not) must have seed length at least $\Omega(\log(n/\varepsilon))$.

1.5 Beyond PRGs: Hitting set generators and more

For the sake of context, in this section we briefly describe some relaxations of the PRG definition. The main motivation behind studying these relaxations is that constructing PRGs is challenging. These “generalized PRGs” are sometimes easier to construct, and yet they suffice for some (but not all) of the applications of PRGs. We only give a short overview of these concepts, since our main focus is true PRGs.

The most well-known “generalized PRG” concept is a *hitting set generator* (HSG).

Definition 1.5.1 (HSGs). *Suppose \mathcal{F} is a class of functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$. An ε -HSG for \mathcal{F} is a function $G: \{0, 1\}^s \rightarrow \{0, 1\}^n$ such that for every $f \in \mathcal{F}$, if $\mathbb{E}[f] \geq \varepsilon$, then there exists some x such that $f(G(x)) = 1$.*

An HSG is a “one-sided PRG.” HSGs have been studied since the 1980s [AKS87] if not earlier. HSGs can be used to derandomize algorithms that have one-sided error, simply by trying all seeds. In some contexts, HSGs can also be used (in nontrivial ways) to derandomize algorithms that have two-sided error [ACR98; ACRT99; BF99; GVW11; CH22].

A few years ago, Braverman, Cohen, and Garg introduced a different generalization of PRGs, called *weighted PRGs* (WPRGs) [BCG20].³

Definition 1.5.2 (WPRG). *Suppose \mathcal{F} is a class of functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$. An ε -WPRG for \mathcal{F} is a pair (G, ρ) , where $G: \{0, 1\}^s \rightarrow \{0, 1\}^n$ and $\rho: \{0, 1\}^s \rightarrow \mathbb{R}$, such that for every $f \in \mathcal{F}$, we have*

$$\left| \mathbb{E}_{U \sim U_s} [f(G(U)) \cdot \rho(U)] - \mathbb{E}[f] \right| \leq \varepsilon.$$

Thus, WPRGs generalize PRGs because we consider sparse *linear* combinations of the outputs of f rather than sparse *convex* combinations of the outputs of f . Several recent works have exploited this extra flexibility to construct WPRGs with better parameters than known PRGs [BCG20; CL20; CDRST21; PV21b; Hoz21].

Yet another generalization of PRGs is the concept of a *deterministic sampler*.

Definition 1.5.3 (Deterministic sampler). *Suppose \mathcal{F} is a class of functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$. An ε -deterministic sampler for \mathcal{F} is a deterministic oracle algorithm A that makes queries to a function $f \in \mathcal{F}$ and outputs a number $A^f \in \mathbb{R}$ such that $|A^f - \mathbb{E}[f]| \leq \varepsilon$.*

The deterministic sampler model isolates a key feature of PRGs, which is that they are useful even if we merely have black-box access to the function f . Deterministic samplers have been discussed (by name) in a couple of recent works [CH22; PV22]. Several older algorithms can also be understood as deterministic samplers [ACR98; ACRT99; ISW99; BF99; GVW11].

²For a counterexample, see the work of Hoza, Pyne, and Vadhan [HPV21].

³In Braverman, Cohen, and Garg’s original paper [BCG20], they speak of “pseudorandom pseudo-distributions.” The “weighted PRG” terminology was introduced later, by Cohen, Doron, Renard, Sberlo, and Ta-Shma [CDRST21].

One can show that these four concepts form a hierarchy:

$$\text{PRG} \implies \text{WPRG} \implies \text{deterministic sampler} \implies \text{HSG}.$$

Thus, PRGs (our focus in this text) are the most desirable of the four.

Chapter 2

Limited Independence and Small-Bias Generators

In this chapter, we study “ k -wise uniform” generators, “small-bias” generators, and simple combinations thereof. What these PRG constructions have in common is that they are closely related to *error correcting codes*. Prior knowledge of coding theory is not necessary to understand the PRGs. The constructions of these PRGs are fairly elementary, but we emphasize that the analyses are interesting and not always trivial. We will build up to showing that these simple PRGs can fool moderately powerful classes of functions, such as bounded-depth circuits and low-degree polynomials over \mathbb{F}_2 .

2.1 Limited independence

2.1.1 Pairwise uniform bits

For our first PRG, let us fool the first nontrivial case of *juntas*.

Definition 2.1.1 (Juntas). *A function f on $\{0, 1\}^n$ is a k -junta if f only depends on at most k variables, i.e.,*

$$f(x) = g(x_{i_1}, \dots, x_{i_k})$$

for some indices $i_1, \dots, i_k \in [n]$ and some function g .

Fooling 1-juntas is trivial. Indeed, let $G: \{0, 1\} \rightarrow \{0, 1\}^n$ be the PRG with seed length 1 given by $G(b) = (b, b, b, \dots, b)$. Then G perfectly fools every 1-junta, because for any i , the i -th bit in the output of the PRG is a uniform bit.

Let us consider the case of 2-juntas. Fooling one specific 2-junta, such as the function $f(x) = x_7 \wedge x_{13}$, is trivial: using a 2-bit seed, we can sample $X_7, X_{13} \in \{0, 1\}$ uniformly and independently at random and set $X_i = 0$ for all $i \notin \{7, 13\}$. The challenge is to construct a single PRG that fools all 2-juntas simultaneously. In other words, the challenge is that when we design the PRG, we don't know in advance which two bits are relevant.

Theorem 2.1.2 (Pairwise uniform bits). *For every $n \in \mathbb{N}$, there is an explicit PRG that perfectly fools 2-juntas on n bits with seed length $\lfloor \log n \rfloor + 1$.*

A distribution X that perfectly fools 2-juntas is also called a *pairwise uniform* distribution, because every two bits of X are uniform over $\{0, 1\}^2$. In practice, people often use the alternative

phrase “pairwise independent.” This practice is a little sloppy, because it doesn’t clarify the marginal distributions of the individual coordinates of X .

A generator for $n = 3$ was described in [Section 1.1.3](#). The solution for larger n is a natural generalization.

Proof of [Theorem 2.1.2](#). Let $s = \lceil \log n \rceil + 1$, and let I_1, \dots, I_n be distinct nonempty subsets of $[s]$. The PRG $G: \{0, 1\}^s \rightarrow \{0, 1\}^n$ is given by

$$G(y) = \left(\bigoplus_{i \in I_1} y_i, \dots, \bigoplus_{i \in I_n} y_i \right). \quad (2.1)$$

To prove that this works, consider sampling $Y \in \{0, 1\}^s$ uniformly at random, and let $j, k \in [n]$ be distinct. Define

$$Z = \bigoplus_{i \in I_j \cap I_k} Y_i, \quad A = \bigoplus_{i \in I_j \setminus I_k} Y_i, \quad B = \bigoplus_{i \in I_k \setminus I_j} Y_i.$$

Then $(G(Y)_j, G(Y)_k) = (Z \oplus A, Z \oplus B)$. Furthermore, Z , A , and B are mutually independent random variables, and A and B are uniformly distributed. Therefore, $(G(Y)_j, G(Y)_k)$ is uniformly distributed over $\{0, 1\}^2$. Therefore, for any function f that only depends on x_j and x_k , the random variables $f(G(U_s))$ and $f(U_n)$ are identically distributed. \square

The seed length in [Theorem 2.1.2](#) is precisely optimal [[ABI86](#); [CGHFRS85](#)], i.e., every pairwise uniform generator has a seed length of at least $\lceil \log n \rceil + 1$.

2.1.2 k -wise uniform bits

For our next PRG, let us fool the class of k -juntas for any k , i.e., we will construct a *k -wise uniform* distribution.

Theorem 2.1.3 (*k -wise uniform bits*). *For every $n, k \in \mathbb{N}$, there is an explicit PRG that perfectly fools k -juntas on n bits with seed length $O(k \log n)$.*

Proof. Let \mathbb{F}_q be a finite field with at least n elements. Let \mathcal{P} be the set of univariate polynomials over \mathbb{F}_q of degrees less than k . Let $z_1, \dots, z_k \in \mathbb{F}_q$ be distinct. In preparation for defining the PRG, define $H: \mathcal{P} \rightarrow \mathbb{F}_q^k$ by

$$H(p) = (p(z_1), \dots, p(z_k)).$$

The function H is injective, because if $H(p) = H(p')$, then $p - p'$ is a polynomial with at least k zeroes of degree less than k , hence $p = p'$. Furthermore, $|\mathcal{P}| = |\mathbb{F}_q^k| = q^k$, since a polynomial $p \in \mathcal{P}$ can be specified by k coefficients from \mathbb{F}_q . Therefore, H is bijective, and hence if $P \in \mathcal{P}$ is sampled uniformly at random, $H(P)$ is a uniform random vector.

Now let $z_1, \dots, z_n \in \mathbb{F}_q$ be distinct, and define $G: \mathcal{P} \rightarrow \mathbb{F}_q^n$ by

$$G(p) = (p(z_1), \dots, p(z_n)).$$

By the above analysis, when $P \in \mathcal{P}$ is sampled uniformly at random, any k coordinates of $G(P)$ are independent and uniform random.

All that remains is to bridge the gap between field elements and bits. Let q be a power of two, so that field elements can be naturally encoded as bitstrings. The seed of our PRG describes a polynomial $p \in \mathcal{P}$ by giving the encodings of its k coefficients; this requires $k \log q = k \cdot \lceil \log n \rceil$ bits if we pick q to be the smallest power of two that is at least n . The output of our PRG is the sequence of first bits of the encodings of the coordinates of $G(p)$. \square




Figure 2.1: A depth-3 decision tree. Note that the function it computes depends on all 6 variables.

The seed length in [Theorem 2.1.3](#) is optimal up to a constant factor for moderate values of k . More precisely, the optimal seed length is $\Theta(k \cdot \log(n/k))$ [[CGHFRS85](#); [ABI86](#); [CL21](#)], which is a slight improvement over [Theorem 2.1.3](#) when $k \geq n^{1-o(1)}$. Even when k is small, the constant factor in the seed length of [Theorem 2.1.3](#) can be improved by roughly a factor of two [[ABI86](#)].

2.1.3 Perfectly fooling shallow decision trees

Next, let us fool shallow *decision trees*, which generalize juntas.

Definition 2.1.4 (Decision trees). *A decision tree over $\{0, 1\}^n$ is a tree, where each internal node is labeled with a variable x_i and has two children, the two edges leading from an internal node to its children are labeled 0 and 1, and each leaf is labeled with an output value (0 or 1). A decision tree computes a function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ by walking from root to leaf according to the values of the variables queried. (See [Figure 2.1](#).)*

Every k -junta can be computed by a depth- k decision tree. To fool decision trees, rather than constructing a new PRG from scratch, we'll show that every PRG for k -juntas automatically fools depth- k decision trees – even though such a tree might compute a function that depends on far more than k variables. This is a common pattern in PRG design: first one designs a PRG for a relatively simple class of functions, and then one proves that such a PRG automatically fools a more sophisticated class of functions.

Proposition 2.1.5 (Perfect PRGs for shallow decision trees). *Let $n, k \in \mathbb{N}$ and let X be a k -wise uniform distribution over $\{0, 1\}^n$. Then X perfectly fools depth- k decision trees. Consequently, there is an explicit PRG with seed length $O(k \log n)$ that perfectly fools depth- k decision trees on n variables.*

Proof. Let f be a depth- k decision tree. Let A be the set of accepting leaves of f , i.e., leaves that are labeled 1. For each leaf $u \in A$, define $f_u: \{0, 1\}^n \rightarrow \{0, 1\}$ by letting $f_u(x) = 1$ if and only if f arrives at u when it reads x . Note that f_u is a k -junta, because its value only depends on the variables queried on the path from the root to u . Furthermore, we can express f as

$$f(x) = \sum_{u \in A} f_u(x).$$

Therefore, by linearity of expectation,

$$\mathbb{E}[f(X)] = \mathbb{E}\left[\sum_{u \in A} f_u(X)\right] = \sum_{u \in A} \mathbb{E}[f_u(X)] = \sum_{u \in A} \mathbb{E}[f_u] = \mathbb{E}\left[\sum_{u \in A} f_u\right] = \mathbb{E}[f]. \quad \square$$

The simple technique in the proof above is quite valuable. Let us abstract it out and generalize it to the case of imperfect PRGs.

Lemma 2.1.6 (Triangle Inequality for PRG Errors). *Let $f_1, \dots, f_k: \{0, 1\}^n \rightarrow \mathbb{R}$ be functions, let $\lambda_0, \dots, \lambda_k \in \mathbb{R}$, and let $f(x) = \lambda_0 + \sum_{i=1}^k \lambda_i \cdot f_i(x)$. Let X be a distribution over $\{0, 1\}^n$, and assume that X fools f_i with error ε_i for each i . Then X fools f with error ε , where*

$$\varepsilon = \sum_{i=1}^k |\lambda_i| \cdot \varepsilon_i.$$

Proof.

$$\begin{aligned} |\mathbb{E}[f(X)] - \mathbb{E}[f]| &= \left| \sum_{i=1}^k \lambda_i \cdot \mathbb{E}[f_i(X)] - \sum_{i=1}^k \lambda_i \cdot \mathbb{E}[f_i] \right| && \text{(Linearity of expectation)} \\ &\leq \sum_{i=1}^k |\lambda_i| \cdot |\mathbb{E}[f_i(X)] - \mathbb{E}[f_i]| && \text{(Standard triangle inequality)} \\ &\leq \sum_{i=1}^k |\lambda_i| \cdot \varepsilon_i && (X \text{ fools } f_i \text{ with error } \varepsilon_i.) \quad \square \end{aligned}$$

2.1.4 Connection with coding theory: Dual codes

For readers with a background in coding theory, the constructions of pairwise and k -wise uniform generators might have felt familiar. Indeed, the constructions are closely related to the Hadamard code and the Reed-Solomon code, respectively. For the sake of those readers who have some familiarity with coding theory, we will now describe a general elegant characterization of exactly which linear codes induce k -wise uniform distributions. Recall that a linear code over \mathbb{F}_2^n is a subspace $C \subseteq \mathbb{F}_2^n$, and its *dual code* is defined as

$$C^\perp = \{x \in \mathbb{F}_2^n : \forall y \in C, \langle x, y \rangle = 0\},$$

where $\langle \cdot, \cdot \rangle$ is the standard dot product over \mathbb{F}_2^n . The minimum distance of a code is the smallest Hamming distance between two distinct codewords. For a linear code C , this distance coincides with the smallest Hamming weight among all nonzero codewords of C .

Proposition 2.1.7 (Connection between k -wise uniformity and coding theory). *Let $C \subseteq \mathbb{F}_2^n$ be a linear subspace, and sample X uniformly at random from C . Then X is k -wise uniform if and only if C^\perp has minimum distance at least $k + 1$.*

Proof. First, suppose X is k -wise uniform. Let $x \in \mathbb{F}_2^n$ be a nonzero vector with Hamming weight at most k . Then $\langle x, X \rangle$ is a uniform random bit, so there is certainly some $y \in C$ such that $\langle x, y \rangle = 1$. Therefore, $x \notin C^\perp$. Since C^\perp is a subspace, it follows that C^\perp has minimum distance at least $k + 1$.

Conversely, suppose C^\perp has minimum distance at least $k + 1$, and consider any k distinct indices $i_1, \dots, i_k \in [n]$. Let $s = \dim(C)$, and let $M \in \mathbb{F}_2^{n \times s}$ be a matrix with image C and rows M_1, \dots, M_n .

Let $x \in \mathbb{F}_2^n$ be an arbitrary nonzero vector supported on the indices i_1, \dots, i_k . Then x has Hamming weight at most k , so $x \notin C^\perp$, i.e., there is some $z \in \mathbb{F}_2^s$ such that

$$0 \neq \langle x, Mz \rangle = \sum_{i=1}^n x_i \cdot \langle M_i, z \rangle = \left\langle \sum_{i=1}^n x_i M_i, z \right\rangle.$$

Therefore, $\sum_{i=1}^n x_i M_i \neq 0$. Since x was arbitrary, this shows that M_{i_1}, \dots, M_{i_k} are linearly independent. Define

$$M' = \begin{bmatrix} M_{i_1} \\ M_{i_2} \\ \vdots \\ M_{i_k} \end{bmatrix}.$$

Since row rank is equal to column rank, there are k linearly independent columns of M' with indices in some set $J = \{j_1, \dots, j_k\}$. Therefore, when $z \in \mathbb{F}_2^s$ is chosen uniformly at random, $M'z$ is a uniform random element of \mathbb{F}_2^k . To see this, note that for any $x \in \mathbb{F}_2^k$ and any fixing of all z_j with $j \notin J$, there is a unique choice of $z_j \in \mathbb{F}_2^k$ for which $M'z = x$. It follows that X is k -wise uniform. \square

[Proposition 2.1.7](#) provides a generic recipe for constructing k -wise uniform distributions from error correcting codes. Recall that the *redundancy* of a code is the difference between the block length and the message length. A binary linear code with block length n , minimum distance $k + 1$, and redundancy s induces a k -wise uniform generator with seed length s and output length n .

2.2 Small-bias distributions

2.2.1 Fooling parities of variables

For our first imperfect PRG, let us fool *parity functions*.

Definition 2.2.1 (Parity functions). *A parity function is a function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ of the form $f(x) = \bigoplus_{i \in S} x_i$ for some set $S \subseteq [n]$.*

Equivalently, we can think of f as a map $\mathbb{F}_2^n \rightarrow \mathbb{F}_2$. Then f is a parity function if and only if $f(x) = \langle a, x \rangle$ for some fixed vector a , where $\langle \cdot, \cdot \rangle$ is the usual inner product, i.e., $\langle a, x \rangle = \sum_{i=1}^n a_i \cdot x_i$. Sometimes it is more convenient to work with $\{\pm 1\}$ -valued functions, in which case parity functions become *character functions*.

Definition 2.2.2 (Character functions). *Let $n \in \mathbb{N}$, and let $S \subseteq [n]$. The character function of S , denoted $\chi_S: \{0, 1\}^n \rightarrow \{\pm 1\}$, is defined by*

$$\chi_S(x) = \prod_{i \in S} (-1)^{x_i}.$$

Note that $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is a parity function if and only if $(-1)^f$ is a character function. Since $(-1)^f = 1 - 2f$, it follows that fooling character functions with error ε is equivalent to fooling parity functions with error $\varepsilon/2$.

Definition 2.2.3 (Bias). *An ε -biased distribution over $\{0, 1\}^n$ is a distribution that ε -fools character functions. Equivalently, a distribution is ε -biased if it $(\varepsilon/2)$ -fools parity functions. An ε -biased generator is an ε -PRG for character functions.*

Theorem 2.2.4 (Small-bias generators [NN93; Per90]). *For every $n \in \mathbb{N}, \varepsilon > 0$, there is an explicit ε -biased generator with output length n and seed length $O(\log(n/\varepsilon))$.*

Remark 2.2.5. *Ideally, we want to design PRGs for interesting and powerful models of computation. The reader might feel that “parity functions” is hardly a “model of computation” at all, and the utility of ε -biased generators is unclear. However, we will see later that any distribution that fools parity functions with sufficiently low error also fools many more interesting models. Furthermore, ε -biased generators are building blocks in many more powerful PRGs.*

Theorem 2.2.4 was proved by Naor and Naor [NN93] and independently by Peralta [Per90]; we’ll present a simpler construction due to Alon, Goldreich, Håstad, and Peralta [AGHP92].

Proof of Theorem 2.2.4. Let $q = n/\varepsilon$, and assume without loss of generality that q is a power of two. As vector spaces over \mathbb{F}_2 , identify \mathbb{F}_q with $\mathbb{F}_2^{\log q}$. Our PRG $G: \mathbb{F}_q \times \mathbb{F}_q \rightarrow \{0, 1\}^n$ is defined by

$$(G(y, z))_i = \langle y, z^i \rangle.$$

To prove that this works, let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be a nonzero parity function, say $f(x) = \bigoplus_{i \in S} x_i$. Then doing arithmetic in \mathbb{F}_2 ,

$$f(G(y, z)) = \sum_{i \in S} \langle y, z^i \rangle = \left\langle y, \sum_{i \in S} z^i \right\rangle.$$

Define $g(z) = \sum_{i \in S} z^i$. Then g is a nonzero polynomial in $\mathbb{F}_q[z]$ of degree at most n , and $f(G(y, z)) = \langle y, g(z) \rangle$. When z is a root of g , obviously $f(G(y, z)) = 0$. On the other hand, when z is not a root of g , if we sample $Y \in \mathbb{F}_q$ uniformly at random, $f(G(Y, z))$ is a uniform random bit. Therefore, when we sample $Y, Z \in \mathbb{F}_q$ independently and uniformly at random,

$$\mathbb{E}_{Y, Z} [f(G(Y, Z))] = \frac{1}{2} \cdot \Pr_Z [g(Z) \neq 0] \in \left[\frac{1}{2} - \frac{n}{2q}, \frac{1}{2} \right].$$

Since $\mathbb{E}[f] = \frac{1}{2}$, our PRG G fools parity functions with error $n/(2q) = \varepsilon/2$, and hence it fools character functions with error ε . \square

2.2.2 A better seed length for parities of few variables

The seed length $O(\log(n/\varepsilon))$ in **Theorem 2.2.4** is asymptotically optimal [AGHP92]. However, we can achieve a better seed length for parities of *just a few variables*, i.e., functions that are simultaneously parity functions and juntas.

Definition 2.2.6 (k -wise ε -bias). *Let X be a distribution over $\{0, 1\}^n$. We say that X is k -wise ε -biased if it ε -fools every character function χ_S for which $|S| \leq k$. Similarly, a k -wise ε -biased generator is an ε -PRG for character functions χ_S that satisfy $|S| \leq k$.*

Theorem 2.2.7 (k -wise ε -biased generators [NN93]). *For every $n, k \in \mathbb{N}$ and every $\varepsilon > 0$, there is an explicit k -wise ε -biased generator with output length n and seed length $O(\log(k/\varepsilon) + \log \log n)$.*

Proof. Let $G: \{0, 1\}^s \rightarrow \{0, 1\}^n$ be a k -wise uniform generator that is also a linear transformation when we think of it as a map between vector spaces, $G: \mathbb{F}_2^s \rightarrow \mathbb{F}_2^n$. Let Y be an ε -biased distribution over $\{0, 1\}^s$. We will show that $G(Y)$ fools parities of at most k bits. Indeed, let $f(x) = \sum_{i \in S} x_i$,

where $x \in \mathbb{F}_2^n$ and $|S| \leq k$. Let $M \in \mathbb{F}_2^{n \times s}$ be the matrix representation of G , with rows $M_1, \dots, M_n \in \mathbb{F}_2^s$. Then for any $y \in \mathbb{F}_2^s$,

$$f(G(y)) = \sum_{i \in S} \langle M_i, y \rangle = \sum_{i \in S} \sum_{j=1}^s M_{ij} y_j = \sum_{j=1}^s \left(\sum_{i \in S} M_{ij} \right) y_j.$$

This is a parity function of the variables y_1, \dots, y_s . Therefore, since Y is ε -biased, $|\mathbb{E}[f(G(Y))] - \mathbb{E}[f(G(U))]| \leq \varepsilon/2$. Furthermore, since G is k -wise uniform and f is a k -junta, $\mathbb{E}[f(G(U))] = \mathbb{E}[f]$. Therefore, $G(Y)$ is k -wise ε -biased. To achieve the promised seed length, we can plug in the constructions of [Theorems 2.1.3](#) and [2.2.4](#) for G and Y respectively. \square

Once again, the seed length of [Theorem 2.2.7](#) is optimal up to constant factors.

2.2.3 Connection with coding theory: Nearly balanced codes

In [Section 2.1.4](#), we saw a connection between k -wise uniform distributions and error correcting codes that “explains” our constructions of k -wise uniform generators ([Theorems 2.1.2](#) and [2.1.3](#)). Now we discuss a similar connection between ε -biased distributions and error correcting codes.

Suppose $C \subseteq \mathbb{F}_2^m$ is a linear code. It is generally desirable for C to have a large minimum weight. Small-biased distributions are equivalent to codes C that also have a small *maximum* weight. Specifically, we say that C is ε -balanced if every nonzero $x \in C$ has relative Hamming weight $\frac{1}{2} \pm \varepsilon$.

Proposition 2.2.8 (Nearly balanced code \iff small-bias distribution). *Let $M \in \mathbb{F}_2^{m \times n}$ be a linear transformation, and let C be the image of M , i.e., $C = \{Ma : a \in \{0, 1\}^n\}$. Sample X uniformly at random from the rows of M , so $X \in \{0, 1\}^n$. Then C is ε -balanced if and only if X is (2ε) -biased.*

Proof. For any nonzero “message” $a \in \{0, 1\}^n$, the relative Hamming weight of Ma is the fraction of rows M_i of M such that $\langle a, M_i \rangle = 1$, i.e., $\Pr[\langle a, X \rangle = 1]$. \square

A nearly balanced code that stretches an n -bit message to an m -bit codeword corresponds to a small-bias generator that stretches a $(\log m)$ -bit seed to an n -bit pseudorandom string. In both problems, it is desirable to minimize m . A natural way to construct a nearly balanced code is to concatenate the Hadamard code with the Reed-Solomon code. Through [Proposition 2.2.8](#), that gives an explicit ε -biased generator similar to the PRG we constructed to prove [Theorem 2.2.4](#). The two constructions are not quite identical. Both have seed length $O(\log(n/\varepsilon))$, so the coding-theory perspective gives an alternative proof of [Theorem 2.2.4](#).

Because of the connection between small-bias distributions and nearly balanced codes, even constant-factor improvements in the seed length of small-bias generators are interesting. Note that a constant factor in the seed length translates to a constant factor in the *exponent* of the codeword length! The seed length in [Theorem 2.2.4](#) is $2 \log(n/\varepsilon) + O(1)$. For moderate ε , the best small-bias generator is a construction by Ta-Shma [[TS17](#)] with seed length¹

$$\log n + 2 \log(1/\varepsilon) + \tilde{O}(\log^{2/3}(1/\varepsilon)).$$

This seed length is extremely close to the nonconstructive bound of $\log n + 2 \log(1/\varepsilon) + O(1)$ ([Proposition 1.3.1](#)), as well as to the lower bound of $\log n + 2 \log(1/\varepsilon) - \log \log(1/\varepsilon) - O(1)$ [[AGHP92](#)]. Ta-Shma’s seed length translates to an ε -balanced code that stretches messages of length n to codewords of length $n/\varepsilon^{2+o(1)}$.

¹Here, we are ignoring rounding issues. That is, the domain size S of Ta-Shma’s generator is not necessarily a power of two, and when we say “seed length” we simply mean $\log_2 S$.

Open Problem 2.2.9 (Optimal small-bias generators up to an *additive* constant in the seed length). *Construct an explicit ε -biased generator with seed length $\log n + 2 \log(1/\varepsilon) + O(1)$, and hence an explicit ε -balanced code that stretches messages of length n to codewords of length $O(n/\varepsilon^2)$.*

2.3 Analysis technique: Fourier L_1 bounds

2.3.1 Basic Fourier analysis

PRGs for character functions (i.e., small-bias distributions) are especially important because character functions are the basic “building blocks” out of which all other functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$ can be assembled.

Proposition 2.3.1 (The Fourier expansion). *Every function $f: \{0, 1\}^n \rightarrow \mathbb{R}$ can be uniquely written as a linear combination of characters, i.e.,*

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \cdot \chi_S(x), \quad (2.2)$$

where $\widehat{f}(S) \in \mathbb{R}$.

Proof. The space of all functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$ is a vector space, isomorphic to \mathbb{R}^{2^n} . Define an inner product on this space by

$$\langle f, g \rangle = \mathbb{E}_{U \sim U_n} [f(U) \cdot g(U)].$$

With respect to this inner product, the 2^n character functions are orthonormal. Therefore, they form a basis. \square

The decomposition of Equation (2.2) is called the *Fourier expansion* of f , and the numbers $\widehat{f}(S)$ are called the *Fourier coefficients* of f . The Fourier expansion of f can reveal important information about f . For example, by linearity of expectation,

$$\mathbb{E}[f] = \sum_{S \subseteq [n]} \widehat{f}(S) \cdot \mathbb{E}[\chi_S] = \widehat{f}(\emptyset). \quad (2.3)$$

2.3.2 Almost k -wise uniform bits

Let us use Fourier analysis to obtain another PRG for k -juntas. For moderate error, its seed length is superior to that of the k -wise uniform generator that we saw before (Theorem 2.1.3). The following theorem is a form of “Vazirani’s XOR Lemma.”

Theorem 2.3.2 (Almost k -wise uniform generator [NN93]). *If X is a k -wise δ -biased distribution over $\{0, 1\}^n$, then X fools $[-1, 1]$ -valued k -juntas with error $\delta \cdot 2^{k/2}$. Consequently, for every $k, n \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for $[-1, 1]$ -valued k -juntas with seed length $O(k + \log(1/\varepsilon) + \log \log n)$.*

A distribution X that fools all $\{0, 1\}$ -valued k -juntas with error ε is also called an ε -almost k -wise uniform distribution.² An equivalent condition is that every k coordinates of X are ε -close to U_k in total variation distance. In practice, people often use the alternative phrase “ ε -almost k -wise independent.”

The proof of Theorem 2.3.2 is based on bounding the magnitude of Fourier coefficients.

²Warning: Occasionally, the same “ ε -almost” terminology refers to some other measure of the extent to which X fails to be perfectly k -wise uniform [Aar10].

Definition 2.3.3 (Fourier L_1 norm). Let $f: \{0,1\}^n \rightarrow \mathbb{R}$. The Fourier L_1 norm of f , denoted $L_1(f)$, is the sum of absolute values of Fourier coefficients of f :

$$L_1(f) = \sum_{S \subseteq [n]} |\widehat{f}(S)|.$$

Lemma 2.3.4 (Universal Fourier L_1 bound). For any function $f: \{0,1\}^n \rightarrow [-1,1]$, we have $L_1(f) \leq 2^{n/2}$.

Proof. By the Cauchy-Schwarz inequality, $L_1(f) \leq \sqrt{2^n \cdot \sum_{S \subseteq [n]} \widehat{f}(S)^2}$. Furthermore, for any function $f: \{0,1\}^n \rightarrow \mathbb{R}$,

$$\mathbb{E}[f(U_n)^2] = \langle f, f \rangle = \sum_{S, T \subseteq [n]} \widehat{f}(S) \cdot \widehat{f}(T) \cdot \langle \chi_S, \chi_T \rangle = \sum_{S \subseteq [n]} \widehat{f}(S)^2. \quad (2.4)$$

(Equation (2.4) is called *Parseval's theorem*.) In our case, f is $[-1,1]$ -valued, so $\mathbb{E}[f(U_n)^2] \leq 1$ and hence $\sum_{S \subseteq [n]} \widehat{f}(S)^2 \leq 1$. \square

Lemma 2.3.5 (Fourier L_1 bound \implies fooled by small-bias). Let $f: \{0,1\}^n \rightarrow \mathbb{R}$. If X is ε -biased, then X fools f with error $\varepsilon \cdot L_1(f)$.

Proof. This is a special case of the Triangle Inequality for PRG Errors (Lemma 2.1.6). \square

Proof of Theorem 2.3.2. Let $f: \{0,1\}^n \rightarrow [-1,1]$ be a k -junta, i.e., $f(x) = g(x_{i_1}, \dots, x_{i_k})$ for some function $g: \{0,1\}^k \rightarrow [-1,1]$, where i_1, \dots, i_k are distinct. By Lemmas 2.3.4 and 2.3.5, the distribution $(X_{i_1}, \dots, X_{i_k})$ fools g with error $\delta \cdot 2^{k/2}$, and hence X fools f with the same error. The final seed length follows from Theorem 2.2.7 by choosing $\delta = \varepsilon \cdot 2^{-k/2}$. \square

2.3.3 Fooling bounded-size decision trees

Recall that in Section 2.1.3, we showed that k -wise uniform generators, with seed length $O(k \log n)$, perfectly fool depth- k decision trees. As another application of Fourier L_1 bounds, let's design another PRG for bounded-depth decision trees with a better seed length (although this time the error will be nonzero). More generally, we will consider decision trees of unbounded depth but bounded *size*. The size of a decision tree is the number of leaves. (See Figure 2.2.)

Proposition 2.3.6 (PRG for bounded-size decision trees). If X is a δ -biased distribution over $\{0,1\}^n$, then X fools size- m decision trees with error $m\delta$. Consequently, for every $n, m \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for size- m decision trees on n input bits with seed length $O(\log(mn/\varepsilon))$.

Note that Proposition 2.3.6 implies a PRG for depth- k decision trees with seed length $O(k + \log(n/\varepsilon))$, because a depth- k decision tree always has size at most 2^k . The proof of Proposition 2.3.6 is similar to the construction of almost k -wise uniform generators: we will bound the Fourier L_1 norm of size- m decision trees. We start with the special case of *conjunctions of literals*.

Proposition 2.3.7 (Fourier L_1 bound for conjunctions of literals). Suppose $f: \{0,1\}^n \rightarrow \{0,1\}$ is a conjunction of literals, i.e.,

$$f(x) = \bigwedge_{i \in S} \ell_i$$

where $S \subseteq [n]$ and each ℓ_i is either x_i or $\neg x_i$. Then $L_1(f) = 1$.




Figure 2.2: A decision tree computing the OR function on n bits. Note that the *size* of this decision tree is $n + 1$, which is relatively low, whereas the *depth* of this decision tree is maximal, which is unavoidable for the OR function.

Proof. There is a convenient formula for the Fourier coefficients of any function f :

$$\mathbb{E}_{U \sim U_n} [f(U) \cdot \chi_S(U)] = \langle f, \chi_S \rangle = \sum_{T \subseteq [n]} \widehat{f}(T) \cdot \langle \chi_T, \chi_S \rangle = \widehat{f}(S). \quad (2.5)$$

In the case $f = \text{NOR}_n$ (the n -input NOR function), we get

$$\widehat{\text{NOR}_n}(S) = \mathbb{E}_{U \sim U_n} [\text{NOR}_n(U) \cdot \chi_S(U)] = 2^{-n}.$$

Therefore, $L_1(\text{NOR}_n) = 2^n \cdot 2^{-n} = 1$. More generally, consider any conjunction of literals f . Without loss of generality, we may assume that all n variables appear in f . Consequently, there is some string $a \in \{0, 1\}^n$ such that $f(x) = \text{NOR}_n(x + a)$, where $+$ is the bitwise XOR operation. Therefore, by Equation (2.5), for each $S \subseteq [n]$,

$$\widehat{f}(S) = \mathbb{E}_U [\text{NOR}_n(U + a) \cdot \chi_S(U)] = \mathbb{E}_U [\text{NOR}_n(U) \cdot \chi_S(U + a)] = \chi_S(a) \cdot \widehat{\text{NOR}_n}(S) = \pm 2^{-n}.$$

(In general, negating variables can only change the signs of Fourier coefficients, not the absolute values.) Therefore, $L_1(f) = 2^n \cdot 2^{-n} = 1$. \square

Corollary 2.3.8 (Fourier L_1 bound for decision trees). *If f is a size- m decision tree, then $L_1(f) \leq m$.*

Proof. One can verify that the Fourier $L_1(f)$ norm truly is a norm, i.e., $L_1(f + g) \leq L_1(f) + L_1(g)$ and $L_1(\lambda f) = |\lambda| \cdot L_1(f)$. Let f be a size- m decision tree. Just like in the proof of Proposition 2.1.5, we can write $f = \sum_{u \in A} f_u$, where A is the set of accepting leaves and $f_u(x)$ indicates whether x leads to u . Each f_u is a conjunction of literals. Therefore, $L_1(f) \leq \sum_{u \in A} L_1(f_u) \leq m$. \square




Figure 2.3: A width-2 length-8 branching program computing the function $f: \{0,1\}^4 \rightarrow \{0,1\}$ defined by $f(x) = 1 \iff |x| = 2$, where $|x|$ denotes Hamming weight. Note that this function cannot be computed by a width-2 *read-once* branching program. This example is derived from work by Borodin, Dolev, Fich, and Paul [BDFP86].

Combining Corollary 2.3.8, Lemma 2.3.5, and Theorem 2.2.4 completes the proof of Proposition 2.3.6. When $m < n$, one can improve the seed length to $O(\log(m/\varepsilon) + \log \log n)$ using m -wise δ -biased generators.

Proposition 2.3.6 extends to the more powerful model of *parity decision trees*, which are decision trees in which each internal node may query an arbitrary parity function of the input [KM93]. The reason is that we can write such a tree f as $f(x) = g(h_1(x), \dots, h_m(x))$ where g is a size- m standard decision tree and h_1, \dots, h_m are parity functions. Consequently,

$$f(x) = \sum_{S \subseteq [n]} \hat{g}(x) \cdot (-1)^{\sum_{i \in S} h_i(x)}.$$

For each fixed S , the function $(-1)^{\sum_{i \in S} h_i(x)}$ is a character function, so it has Fourier L_1 norm 1, and hence $L_1(f) \leq L_1(g) \leq m$.

2.3.4 Fooling width-2 branching programs

For a final application of Fourier L_1 bounds, let us obtain a PRG for *width-2 branching programs* (see Figure 2.3). Branching programs are one of the oldest sequential models of computation. For general size- m branching programs, which model computing with $\log m$ bits of memory, the current best PRG has seed length roughly \sqrt{m} (see Section 5.6). We can do much better in the special case of width-2 branching programs, which model computing with a *single bit of memory* and a clock. The precise definition follows.

Definition 2.3.9 (Bounded-width branching programs). *A width- w length- m branching program f is a directed (multi)graph with $m + 1$ layers V_0, \dots, V_m of w vertices each. For $i \in [m]$, each vertex $v \in V_{i-1}$ is labeled with an index $j_v \in [n]$ and has two outgoing edges labeled 0 and 1 leading to vertices in V_i . There is a designated “start vertex” $v_{\text{start}} \in V_0$ and a designated set of “accepting vertices” $V_{\text{accept}} \subseteq V_m$.³ Given an input $x \in \{0,1\}^n$, the program starts at v_{start} , and in each step, having reached a vertex v , the program queries x_{j_v} and traverses the corresponding outgoing edge. Eventually, the program reaches a vertex $v \in V_m$, and $f(x) = 1 \iff v \in V_{\text{accept}}$.*

The following theorem is attributed to unpublished 1995 work of Saks and Zuckerman (see also the work of Bogdanov, Dvir, Verbin, and Yehudayoff [BDVY13]).

³Note that we do not allow the branching program to halt prior to reaching layer m . This type of program is sometimes referred to as a “strict” width- w program [BDFP86].

Theorem 2.3.10 (PRGs for width-2 branching programs). *If X is a δ -biased distribution over $\{0, 1\}^n$, then X fools width-2 length- m branching programs with error $\delta \cdot (m + 1)/2$. Consequently, for every $n, m \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for width-2 length- m branching programs with seed length $O(\log(mn/\varepsilon))$.*

One can show that every width-2 branching program on n variables can be simulated by a width-2 branching program of length $m = O(n^2)$ [BDFP86], so the seed length in Theorem 2.3.10 actually simplifies to $O(\log(n/\varepsilon))$. When Theorem 2.3.10 is mentioned in the literature, it is sometimes indicated that we should assume that the branching program is read-once [GMRTV12; BDVY13; HZ20], but such an assumption is not necessary. Once again, we will prove Theorem 2.3.10 by proving a Fourier L_1 bound.

Lemma 2.3.11 (Fourier L_1 bound for width-2 branching programs). *If f is a width-2 length- m branching program, then $L_1(f) \leq m/2 + 1/2$.*

Proof. Let $F(x) = (-1)^{f(x)}$. For each vertex v in f , let $f_{\rightarrow v}(x)$ indicate whether $f(x)$ visits v , and let $F_{\rightarrow v}(x) = (-1)^{f_{\rightarrow v}(x)}$. We will prove by induction on m that $L_1(F) \leq m$. For the base case $m = 1$, the function F is a 1-junta, i.e., $F(x) = (-1)^{x_i}$ or $F(x) = (-1)^{1-x_i}$ or $F(x) = 1$ or $F(x) = -1$. In each case, $L_1(F) = 1$. Now, for the inductive step, let $V_{m-1} = \{u, v\}$. Then there exist 1-juntas $\phi_u, \phi_v: \{0, 1\}^n \rightarrow \{\pm 1\}$ such that

$$\begin{aligned} F(x) &= f_{\rightarrow u}(x) \cdot \phi_u(x) + f_{\rightarrow v}(x) \cdot \phi_v(x) \\ &= \left(\frac{1}{2} - \frac{1}{2} \cdot F_{\rightarrow u}(x)\right) \cdot \phi_u(x) + \left(\frac{1}{2} + \frac{1}{2} \cdot F_{\rightarrow u}(x)\right) \cdot \phi_v(x) \\ &= \frac{1}{2} \cdot F_{\rightarrow u}(x) \cdot (\phi_v(x) - \phi_u(x)) + \frac{1}{2} \cdot (\phi_u(x) + \phi_v(x)). \end{aligned}$$

Now, $L_1(\phi_u) = L_1(\phi_v) = 1$, and by induction, $L_1(F_{\rightarrow u}) \leq m - 1$. Therefore,

$$L_1(F) \leq \frac{1}{2} \cdot L_1(F_{\rightarrow u}) \cdot (L_1(\phi_v) + L_1(\phi_u)) + \frac{1}{2} (L_1(\phi_u) + L_1(\phi_v)) \leq m - 1 + 1 = m,$$

completing the induction. Finally, $f(x) = \frac{1}{2} - \frac{1}{2}F(x)$, so $L_1(f) \leq \frac{1}{2} + \frac{m}{2}$. \square

Theorem 2.3.10 follows by combining Lemmas 2.3.5 and 2.3.11 and Theorem 2.2.4. Analogously to the situation with decision trees, when $m < n$, one can improve the seed length to $O(\log(m/\varepsilon) + \log \log n)$ using $(2m)$ -wise δ -biased generators.

2.4 Viola's generator for low-degree \mathbb{F}_2 -polynomials

In Section 2.2 we saw a simple construction of an explicit small-bias generator, i.e., a PRG that fools all \mathbb{F}_2 -linear functions with logarithmic seed length. We've discussed connections between small-bias generators and coding theory and some simple applications of small-bias generators. As a natural generalization, let us construct PRGs for quadratic or higher degree polynomials.

Remark 2.4.1 (Polynomials over \mathbb{F}_2 vs. polynomials over \mathbb{R}). *Over the reals, every degree- d polynomial is perfectly fooled by d -wise uniform generators. However, in this section, we are working over \mathbb{F}_2 . Thus, a low-degree polynomial is a function of the form PARITY \circ AND where the AND gates have low fan-in. In this setting, k -wise uniformity is not a good approach. For instance, the uniform distribution over all strings with even Hamming weight is $(n - 1)$ -wise uniform, and yet it does not even fool degree-1 polynomials (parity functions).*

The problem of designing PRGs for low-degree \mathbb{F}_2 -polynomials seemed to be much harder than constructing small-bias generators or k -wise uniform generators. For a long time, even for constant degree, the best construction known was a PRG by Luby, Veličković, and Wigderson [LVW93] with seed length $2^{O(\sqrt{\log n})}$. Over a decade later, a new line of work [Bog05; Lov09] led to Viola’s elegant proof [Vio09] that simply summing d independent copies of small-bias generators gives a PRG for degree- d polynomials.

Theorem 2.4.2 (PRG for low-degree \mathbb{F}_2 -polynomials [Vio09]). *Let Y_1, \dots, Y_d be independent δ -biased random variables over \mathbb{F}_2^n where $\delta \leq 1/2$. Then $Y_1 + \dots + Y_d$ fools degree- d \mathbb{F}_2 -polynomials with error $4 \cdot (\delta/2)^{1/2^{d-1}}$. Consequently, for every $n, d \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for degree- d \mathbb{F}_2 -polynomials in n variables with seed length $O(d \cdot \log n + d \cdot 2^d \cdot \log(1/\varepsilon))$.*

For context, it is easy to show that a sum of independent small-bias random variables is “more pseudorandom” than a single small-bias random variable in the sense that it has smaller bias (see below). [Theorem 2.4.2](#) says that not only does the sum have smaller bias, it also fools higher-degree polynomials.

Observation 2.4.3 (XORing decreases bias). *Let Y_1, \dots, Y_d be independent δ -biased random variables distributed over \mathbb{F}_2^n . Then $\sum_{i=1}^d Y_i$ is (δ^d) -biased.*

Proof. For every nonempty $S \subseteq [n]$, we have

$$\left| \mathbb{E} \left[\chi_S \left(\sum_{i=1}^d Y_i \right) \right] \right| = \left| \prod_{i=1}^d \mathbb{E}[\chi_S(Y_i)] \right| \leq \delta^d. \quad \square$$

2.4.1 Directional derivatives

The proof of [Theorem 2.4.2](#) is based on the notion of *directional derivatives* over \mathbb{F}_2 , defined below.

Definition 2.4.4 (Directional derivative). *Let $f: \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ and $y \in \mathbb{F}_2^n$. The directional derivative $\partial_y f$ is defined by*

$$\partial_y f(x) = f(x + y) + f(x).$$

If \mathcal{F} is a class of functions $f: \mathbb{F}_2^n \rightarrow \mathbb{F}_2$, we define $\partial \mathcal{F} = \{\partial_y f : f \in \mathcal{F}, y \in \mathbb{F}_2^n\}$.

To fool low-degree polynomials, our strategy will be to show how to convert PRGs for $\partial \mathcal{F}$ into PRGs for \mathcal{F} , where \mathcal{F} is any “reasonable” class. Formally, the only requirement on \mathcal{F} is that is “closed under shifts,” as defined below.

Definition 2.4.5 (Closure under shifts). *For a function f on \mathbb{F}_2^n and a vector $y \in \mathbb{F}_2^n$, we define the shift f^{+y} by the formula $f^{+y}(x) = f(x + y)$. Let \mathcal{F} be a class of functions f on \mathbb{F}_2^n . We say that \mathcal{F} is closed under shifts if for every $f \in \mathcal{F}$ and every $y \in \mathbb{F}_2^n$, we have $f^{+y} \in \mathcal{F}$.*

Lemma 2.4.6 (PRG for $\partial \mathcal{F} \implies$ PRG for \mathcal{F}). *Let \mathcal{F} be a class of functions $f: \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ that is closed under shifts. Suppose W fools $\partial \mathcal{F}$ with error γ , Y is δ -biased, and Y is independent of W . Then $W + Y$ fools \mathcal{F} with error $\sqrt{2\gamma} + \delta/2$.*

In general, $\partial \mathcal{F}$ seems to be “more complicated” than \mathcal{F} itself, so [Lemma 2.4.6](#) might not sound particularly useful. However, $\partial \mathcal{F}$ is “simpler” than \mathcal{F} in one respect, namely degree:

Observation 2.4.7 (Differentiation decreases degree). *Let $d \geq 1$, let $f: \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ be a degree- d polynomial, and let $y \in \mathbb{F}_2^n$. Then $\partial_y f$ is a degree- $(d - 1)$ polynomial.*

Thus, we will be able to prove [Theorem 2.4.2](#) by applying [Lemma 2.4.6](#) inductively.

2.4.2 The XOR of two independent copies of an arbitrary Boolean function

The proof of [Lemma 2.4.6](#) (the reduction from fooling \mathcal{F} to fooling $\partial\mathcal{F}$) relies on the following lemma, which explains how to use small-bias distributions to “recycle” randomness and thereby fool a certain class of functions.

Lemma 2.4.8 (Using small-bias distributions to fool $g(x) \cdot g(y)$). *Let n be an even positive integer, let $g: \{0, 1\}^{n/2} \rightarrow \{\pm 1\}$, and let $f(x, y) = g(x) \cdot g(y)$. Let U and Y be independent, where $U \sim U_{n/2}$ and Y is an ε -biased random variable over $\mathbb{F}_2^{n/2}$. Then $(U, U + Y)$ fools f with error ε .*

Proof. Define $F: \{0, 1\}^{n/2} \rightarrow [-1, 1]$ by

$$F(x) = \mathbb{E}_U[g(U) \cdot g(U + x)].$$

Let $U' \sim U_{n/2}$ be independent of U . Then for any $S \subseteq [n]$, by [Equation \(2.5\)](#),

$$\begin{aligned} \widehat{F}(S) &= \mathbb{E}_{U, U'}[g(U) \cdot g(U + U') \cdot \chi_S(U')] = \mathbb{E}_{U, U'}[g(U) \cdot g(U') \cdot \chi_S(U + U')] = \left(\mathbb{E}_U[g(U) \cdot \chi_S(U)] \right)^2 \\ &= \widehat{g}(S)^2. \end{aligned}$$

Therefore, $L_1(F) = \sum_S \widehat{g}(S)^2 \leq 1$ by Parseval’s theorem. Consequently, Y fools F with error ε by [Lemma 2.3.5](#), and hence $(U, U + Y)$ fools f with error ε . \square

Remark 2.4.9 (Characterizing small bias). *One can show the following converse to [Lemma 2.4.8](#): If $U \sim U_{n/2}$, Y is independent, and $(U, U + Y)$ fools all functions of the form $f(x, y) = g(x) \cdot g(y) \in \{\pm 1\}$ with error ε , then Y is ε -biased. Thus, the condition in [Lemma 2.4.8](#) gives an alternative characterization of small-bias distributions.*

Remark 2.4.10 (Connection to expander graphs). *The seed length for sampling the distribution $(U, U + Y)$ that appears in [Lemma 2.4.8](#) is $n/2 + O(\log(n/\varepsilon))$. [Lemma 2.4.8](#) generalizes to the case that g has bounded variance, $\text{Var}[g] \leq 1$, rather than being $\{\pm 1\}$ -valued. This generalization is closely related to the notion of a spectral expander. In [Section 3.1.1](#), we will discuss spectral expanders in more detail, and in particular we will discuss PRGs for such tests with the improved seed length $n/2 + O(\log(1/\varepsilon))$.*

2.4.3 The reduction from fooling \mathcal{F} to fooling $\partial\mathcal{F}$

To prove [Lemma 2.4.6](#), we will consider two cases based on the extent to which $f \in \mathcal{F}$ is balanced. For a function $f: \mathbb{F}_2^n \rightarrow \mathbb{F}_2$, define

$$\text{imbalance}(f) = \left| \mathbb{E} \left[(-1)^{f(U_n)} \right] \right| = 2 \cdot \left| \mathbb{E}[f] - \frac{1}{2} \right|.$$

(In the literature, this quantity is often referred to as the “bias” of f . We use the term “imbalance” instead to avoid confusion with small-bias distributions.) We begin with the case that f is close to balanced.

Lemma 2.4.11 (Fooling well-balanced functions). *Let \mathcal{F} be a class of functions $f: \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ that is closed under shifts. Suppose W fools $\partial\mathcal{F}$ with error γ , Y is δ -biased, and Y is independent of W . Then $W + Y$ fools each $f \in \mathcal{F}$ with error $\text{imbalance}(f) + \sqrt{\gamma/2} + \delta/2$.*

Proof. First observe that

$$\begin{aligned} |\mathbb{E}[f(W + Y)] - \mathbb{E}[f(U_n)]| &= \frac{1}{2} \cdot \left| \mathbb{E} \left[(-1)^{f(W+Y)} \right] - \mathbb{E} \left[(-1)^{f(U_n)} \right] \right| \\ &\leq \frac{1}{2} \cdot \left| \mathbb{E} \left[(-1)^{f(W+Y)} \right] \right| + \frac{1}{2} \cdot \text{imbalance}(f). \end{aligned}$$

Thus, it suffices to bound $|\mathbb{E}[(-1)^{f(W+Y)}]|$. By Jensen's inequality,

$$\left(\mathbb{E}_{W,Y} \left[(-1)^{f(W+Y)} \right] \right)^2 \leq \mathbb{E}_W \left[\left(\mathbb{E}_Y \left[(-1)^{f(W+Y)} \right] \right)^2 \right] = \mathbb{E}_{W,Y,Y'} \left[(-1)^{f(W+Y)+f(W+Y')} \right],$$

where Y' is an independent copy of Y . For fixed Y , the function $f^{+Y}(x) \stackrel{\text{def}}{=} f(x+Y)$ is in \mathcal{F} because \mathcal{F} is closed under shifts. Furthermore, for fixed Y, Y' , the function $g(x) \stackrel{\text{def}}{=} f(x+Y) + f(x+Y')$ is in $\partial\mathcal{F}$, because $g = \partial_{Y+Y'} f^{+Y}$. Therefore, the assumption on W gives

$$\mathbb{E}_{W,Y,Y'} \left[(-1)^{f(W+Y)+f(W+Y')} \right] \leq \mathbb{E}_{Y,Y',U \sim U_n} \left[(-1)^{f(U+Y)+f(U+Y')} \right] + 2\gamma.$$

Finally, observing that $(U+Y, U+Y')$ is identically distributed to $(U, U+Y+Y')$, we have

$$\mathbb{E}_{Y,Y',U \sim U_n} \left[(-1)^{f(U+Y)+f(U+Y')} \right] = \mathbb{E}_{Y,Y',U \sim U_n} \left[(-1)^{f(U)+f(U+Y+Y')} \right] \leq \text{imbalance}(f)^2 + \delta^2,$$

where the last inequality follows from [Lemma 2.4.8](#) and the fact that $Y+Y'$ is (δ^2) -biased. In summary, we have shown that

$$\begin{aligned} |\mathbb{E}[f(W + Y)] - \mathbb{E}[f(U_n)]| &\leq \frac{1}{2} \cdot \left| \mathbb{E}_{W,Y} \left[(-1)^{f(W+Y)} \right] \right| + \frac{1}{2} \cdot \text{imbalance}(f) \\ &\leq \frac{1}{2} \cdot \sqrt{\text{imbalance}(f)^2 + \delta^2 + 2\gamma} + \frac{1}{2} \cdot \text{imbalance}(f) \\ &\leq \text{imbalance}(f) + \frac{\delta}{2} + \sqrt{\frac{\gamma}{2}}. \quad \square \end{aligned}$$

Now we move on to the case that f is significantly imbalanced. In this case, W alone (rather than $W+Y$) already fools f .

Lemma 2.4.12 (Fooling imbalanced functions). *Let \mathcal{F} be any class of functions $f: \mathbb{F}_2^n \rightarrow \mathbb{F}_2$. Suppose W fools $\partial\mathcal{F}$ with error γ . Then W fools each $f \in \mathcal{F}$ with error $\gamma \cdot \text{imbalance}(f)^{-1}$.*

Proof. Let U and U' be two independent copies of U_n . Then

$$\begin{aligned} \text{imbalance}(f) \cdot \left| \mathbb{E} \left[(-1)^{f(W)} \right] - \mathbb{E} \left[(-1)^{f(U)} \right] \right| &= \left| \mathbb{E} \left[(-1)^{f(W)+f(U)} \right] - \mathbb{E} \left[(-1)^{f(U)+f(U')} \right] \right| \\ &= \left| \mathbb{E} \left[(-1)^{f(W)+f(W+U)} \right] - \mathbb{E} \left[(-1)^{f(U')+f(U'+U)} \right] \right| \\ &\leq 2\gamma, \end{aligned}$$

where the last inequality is due to the fact that for any fixing of U , the function $\partial_U f$ is fooled by W . \square

Now we combine the two cases to complete the proof of the reduction.

Proof of Lemma 2.4.6. For any $f \in \mathcal{F}$ and any fixing of Y , the function $f^{+Y}(x) \stackrel{\text{def}}{=} f(x+Y)$ is in \mathcal{F} , and $\text{imbalance}(f^{+Y}) = \text{imbalance}(f)$. Therefore, Lemma 2.4.12 implies that $W+Y$ fools f with error $\gamma \cdot \text{imbalance}(f)^{-1}$. Combining with Lemma 2.4.11 shows that $W+Y$ fools f with error

$$\min \left\{ \gamma \cdot \text{imbalance}(f)^{-1}, \text{imbalance}(f) + \sqrt{\gamma/2} + \delta/2 \right\} \leq \sqrt{2\gamma} + \delta/2,$$

where the last inequality follows by case analysis based on whether $\text{imbalance}(f) \leq \sqrt{\gamma/2}$. \square

2.4.4 Inductive analysis of low-degree polynomials

Proof of Theorem 2.4.2. By Lemma 2.4.6 and Observation 2.4.7, for every i , the random variable $\sum_{j=1}^i Y_j$ fools degree- i polynomials with error ε_i , where $\varepsilon_1 = \delta/2$ and $\varepsilon_{i+1} = \sqrt{2\varepsilon_i} + \delta/2$. Since $\delta \leq 1/2$, we get $\varepsilon_{i+1} \leq \sqrt{2\varepsilon_i} + \sqrt{\delta/2}/2$. Since $\varepsilon_i \geq \delta/2$, we get $\varepsilon_{i+1} \leq (\sqrt{2} + 1/2) \cdot \sqrt{\varepsilon_i} \leq 2\sqrt{\varepsilon_i}$. It follows that

$$\varepsilon_d \leq 4 \cdot (\delta/2)^{1/2^{d-1}}.$$

The seed length bound follows by choosing $\delta = 2 \cdot (\varepsilon/4)^{2^{d-1}}$ and using the small-bias generator construction of Theorem 2.2.4. \square

When d is constant, the seed length in Theorem 2.4.2 is optimal. However, the generator becomes trivial when $d = \Theta(\log n)$.

Open Problem 2.4.13 (PRGs for logarithmic-degree polynomials). *Design an explicit nontrivial PRG for \mathbb{F}_2 -polynomials of degree $\log n$.*

Open Problem 2.4.13 is closely related to the challenge of proving better *correlation bounds* against polynomials; see Viola’s survey [Vio22].

2.4.5 Application: Width-2 branching programs that read several bits at a time

Studying low-degree polynomials is natural enough from a mathematical perspective, but what about from a computing perspective? The reader might find it strange to think of polynomials as a computational model. However, we will now show that PRGs for low-degree polynomials imply PRGs for other models of a more “computational” nature, which demonstrates the importance of Viola’s PRG. In particular, we can fool *compositions with juntas*, provided that the outer function has bounded Fourier L_1 norm.

Definition 2.4.14 (Compositions with juntas). *Let $f: \{0, 1\}^r \rightarrow \mathbb{R}$. For each $n, d \in \mathbb{N}$, we define $f \circ \text{JUNTA}_{n,d}$ to be the class of all functions $g: \{0, 1\}^n \rightarrow \mathbb{R}$ of the form*

$$g(x) = f(\phi_1(x), \dots, \phi_r(x)),$$

where each ϕ_i is a d -junta on n bits. If \mathcal{F} is a class of functions $f: \{0, 1\}^r \rightarrow \mathbb{R}$, then we define $\mathcal{F} \circ \text{JUNTA}_{n,d} = \bigcup_{f \in \mathcal{F}} f \circ \text{JUNTA}_{n,d}$.

Lemma 2.4.15 (PRGs for compositions with juntas). *Suppose X is a distribution over $\{0, 1\}^n$ that fools degree- d polynomials over \mathbb{F}_2^n with error ε , and let $f: \{0, 1\}^r \rightarrow \mathbb{R}$. Then X fools $f \circ \text{JUNTA}_{n,d}$ with error $2\varepsilon \cdot L_1(f)$.*

Proof. If $g(x) = f(\phi_1(x), \dots, \phi_r(x))$, then by the Fourier expansion of f , we have

$$g(x) = \sum_{S \subseteq [r]} \widehat{f}(S) \cdot (-1)^{\sum_{i \in S} \phi_i(x)}.$$

The summation in the exponent may be performed modulo 2. If each ϕ_i is a d -junta, then each ϕ_i can be computed by a degree- d polynomial over \mathbb{F}_2 , hence $\sum_{i \in S} \phi_i(x) \bmod 2$ is also a degree- d polynomial over \mathbb{F}_2 . Therefore, X fools $(-1)^{\sum_{i \in S} \phi_i(x)}$ with error 2ε . The lemma follows by the Triangle Inequality for PRG Errors. \square

Probably the most interesting example is when we take \mathcal{F} to be the class of width-2 length- m branching programs on $2m$ input bits (see [Definition 2.3.9](#)). Then $\mathcal{F} \circ \text{JUNTA}_{n,d}$ is precisely the class of functions computable by a variant model of width-2 length- m branching programs in which *the program reads d bits at a time*, i.e., each vertex v is labeled by a set of indices $J_v \subseteq [n]$ with $|J_v| = d$ and has 2^d outgoing edges corresponding to the possible values of the input substring x_{J_v} . For this model, [Theorem 2.4.2](#) and [Lemmas 2.3.11](#) and [2.4.15](#) imply a seed length of $O(d \cdot \log n + d \cdot 2^d \cdot \log(m/\varepsilon))$. This was shown by Bogdanov, Dvir, Verbin, and Yehudayoff [[BDVY13](#)]. (They assume that the program is “oblivious” in the sense that $J_u = J_v$ if u and v are in the same layer, but such an assumption is not necessary.)

2.5 Analysis technique: Sandwiching approximators

2.5.1 The sandwiching lemma

Suppose we wish to show that a distribution X fools some class \mathcal{F} . A common approach has two steps:

1. Prove that X fools some “simpler” class $\mathcal{F}_{\text{simp}}$.
2. Prove a “transfer theorem,” saying that every distribution that fools $\mathcal{F}_{\text{simp}}$ also fools \mathcal{F} (possibly with some loss in the error parameter).

The second step requires showing that $\mathcal{F}_{\text{simp}}$ can “simulate” \mathcal{F} in some sense. For example, several times, we have shown that every function in some class of interest in can be written as a *linear combination* of “simpler” functions:

- Every depth- k decision tree can be written as a sum of k -juntas ([Proposition 2.1.5](#)).
- Every Boolean function can be written as a linear combination of parity functions ([Proposition 2.3.1](#)).
- Every width-2 branching program that reads several bits at a time can be written as a linear combination of low-degree polynomials over \mathbb{F}_2 ([Section 2.4.5](#)).

In each case, the Triangle Inequality for PRG Errors gives us our desired transfer theorem. (The final error depends on the magnitude of the coefficients in the linear combination.) In this section, we present a second method for proving a “transfer theorem” stating that every distribution that fools $\mathcal{F}_{\text{simp}}$ also fools \mathcal{F} .

Suppose X is a distribution that fools $\mathcal{F}_{\text{simp}}$, and suppose that $\mathcal{F}_{\text{simp}}$ *approximately* simulates \mathcal{F} in some sense. For example, suppose that for every $f \in \mathcal{F}$, there is an $f' \in \mathcal{F}_{\text{simp}}$ such that $\mathbb{E}[|f - f'|]$ is small. Unfortunately, it does not immediately follow that X fools \mathcal{F} : although f and

f' behave similarly under the uniform distribution, it isn't clear whether they behave similarly under the pseudorandom distribution X . A technique for getting around this issue is to establish a stronger form of approximation called *sandwiching*.

Definition 2.5.1 (Sandwiching). *Let $f, f_\ell, f_u: \{0, 1\}^n \rightarrow \mathbb{R}$. We say that f is δ -sandwiched between f_ℓ and f_u if $f_\ell \leq f \leq f_u$ and $\mathbb{E}[f_u - f_\ell] \leq \delta$. In this case, we refer to f_ℓ and f_u as “sandwichers” or “sandwiching approximators” for f .*

Lemma 2.5.2 (Sandwiching Lemma). *Suppose f is δ -sandwiched between f_ℓ and f_u , and suppose X fools f_ℓ and f_u with error ε . Then X fools f with error $\varepsilon + \delta$.*

Proof.

$$\begin{aligned} \mathbb{E}[f(X)] &\leq \mathbb{E}[f_u(X)] \leq \mathbb{E}[f_u] + \varepsilon \leq \mathbb{E}[f] + \varepsilon + \delta \\ \mathbb{E}[f(X)] &\geq \mathbb{E}[f_\ell(X)] \geq \mathbb{E}[f_\ell] - \varepsilon \geq \mathbb{E}[f] - \varepsilon - \delta. \end{aligned} \quad \square$$

2.5.2 Using k -wise uniform generators to fool size- m decision trees

To illustrate the sandwiching technique, let us return to the decision tree model. Recall that we showed that k -wise uniform generators fool depth- k decision trees (Proposition 2.1.5), and then later we showed that small-bias generators fool size- m decision trees (Proposition 2.3.6). The latter model generalizes the former by taking $m = 2^k$. We now show that k -wise uniform generators also fool bounded-*size* decision trees.

Proposition 2.5.3 (Limited independence fools bounded-size decision trees). *If X is a k -wise uniform distribution, then X fools size- m decision trees with error $m \cdot 2^{-k}$.*

Proof. Let f be a size- m decision tree. Define a depth- k decision tree f_ℓ by starting with f and replacing each internal node at depth exactly k with a leaf labeled 0 (and deleting all of its descendants). Similarly, define f_u by replacing each internal node at depth k with a leaf labeled 1. Let us show that f is δ -sandwiched between f_ℓ and f_u , for $\delta = m \cdot 2^{-k}$.

Clearly $f_\ell \leq f \leq f_u$. For each “new” leaf u of f_ℓ or f_u (i.e., u was not a leaf in f), the probability of reaching u on a uniform random input is precisely 2^{-k} . The number of new leaves is the number of internal nodes of f at depth k , which is at most m . Therefore, by the union bound, $\mathbb{E}[f_u - f_\ell] \leq m \cdot 2^{-k}$.

The Sandwiching Lemma completes the proof, because X fools f_ℓ and f_u with error 0 (see Section 2.1.3). □

Proposition 2.5.3 implies that using k -wise uniform generators, we can ε -fool size- m decision trees using a seed of length $O(\log(m/\varepsilon) \cdot \log n)$. This seed length is *inferior* to the seed length that we obtained previously using small-bias generators, which was $O(\log(mn/\varepsilon))$ (see Proposition 2.3.6). However, sometimes it is useful to understand the effect of specific classes of distributions, such as k -wise uniform distributions, on a given model of computation.

2.5.3 Small-bias distributions fool read-once AC^0

For a more sophisticated example of a sandwiching argument, let us consider “ AC^0 circuits,” i.e., bounded-depth Boolean circuits of unbounded fan-in.




Figure 2.4: A depth-3 read-once \mathbf{AC}^0 formula.

Definition 2.5.4 (\mathbf{AC}^0 circuits). *An \mathbf{AC}^0 circuit is a directed acyclic graph where every input node is labeled by a literal (x_i or $\neg x_i$) or a constant (0 or 1), every internal node (“gate”) is labeled by \wedge or \vee , and there is exactly one output node. The in-degrees (also called fan-ins) of \wedge or \vee gates are not bounded. The size of the circuit is the total number of its \wedge and \vee gates. The depth of the circuit is the length of its longest directed path.*

Traditionally, the expression “ \mathbf{AC}^0 ” refers to the *complexity class* consisting of all languages that can be decided by constant-depth polynomial size families of unbounded-fan-in circuits. As suggested by [Definition 2.5.4](#), we will instead adopt the convenient convention of speaking of “size- m depth- d \mathbf{AC}^0 circuits,” where m is not necessarily $\text{poly}(n)$ and d is not necessarily $O(1)$. That being said, $m = \text{poly}(n)$ and $d = O(1)$ is the parameter regime in which we are most interested.

Later, we will present PRGs for general \mathbf{AC}^0 circuits. (See [Sections 2.6, 4.2, 5.1](#) and [5.3](#).) For now, let us focus on fooling the *read-once* version of \mathbf{AC}^0 , a substantially easier problem. A *read-once \mathbf{AC}^0 formula* is an \mathbf{AC}^0 circuit in which every variable appears at most once and the underlying graph structure is a tree. See [Figure 2.4](#).

Theorem 2.5.5 (Small-bias fools read-once \mathbf{AC}^0). *For every $n, d \in \mathbb{N}$ and $\varepsilon > 0$ with $d \geq 2$, there is a value $\delta = \exp(-\Theta(\log n)^{d-1} \cdot \log(1/\varepsilon))$ such that if X is a δ -biased distribution over $\{0, 1\}^n$, then X fools depth- d read-once \mathbf{AC}^0 formulas with error ε . Consequently, there is an explicit ε -PRG for depth- d read-once \mathbf{AC}^0 formulas with seed length $O(\log n)^{d-1} \cdot \log(1/\varepsilon)$.*

When $d = 2$ (read-once CNFs and DNFs) and ε is constant, the seed length of [Theorem 2.5.5](#) is $O(\log n)$, which is optimal. For larger d or smaller ε , the seed length is not optimal: the optimal seed length would be $O(\log(n/\varepsilon))$, independent of depth (note we always have $d \leq n$). That being said, a benefit of [Theorem 2.5.5](#) is the simplicity of the PRG itself. See [Section 5.5](#) for a discussion of more sophisticated PRGs for read-once \mathbf{AC}^0 with better seed lengths.

The case $d = 2$ of [Theorem 2.5.5](#) was first explicitly stated and proven by De, Etesami, Trevisan, and Tulsiani [[DETT10](#)]. It also readily follows [[LV17](#), Appendix A] from earlier work by Chari, Rohatgi, and Srinivasan [[CRS00](#)]. It seems that the case $d > 2$ does not appear in the literature, but the argument for $d > 2$ is a straightforward generalization of the argument for $d = 2$.

The proof of [Theorem 2.5.5](#) works by repeatedly applying the following lemma.

Lemma 2.5.6 (PRG for depth $d \implies$ PRG for depth $d + 1$). *Suppose a distribution X fools depth- d read-once \mathbf{AC}^0 formulas with error ε , where $d \geq 1$. Then X fools depth- $(d + 1)$ read-once \mathbf{AC}^0 formulas with error $\exp\left(-\Omega\left(\frac{\log(1/\varepsilon)}{\log n}\right)\right)$.*

Proof. Let f be a depth- $(d + 1)$ read-once \mathbf{AC}^0 formula. Assume for now that the output gate of f is \vee , so we can write $f(x) = f_1(x) \vee \cdots \vee f_m(x)$. Define the *weight* of such a formula to be the expected number of terms satisfied on a uniform random input, i.e., $\text{Weight}(f) = \sum_{i=1}^m \mathbb{E}[f_i]$. As a first step, we will show that for every even positive integer k , the distribution X fools f with error

$$\varepsilon \cdot (2m)^k + (e \cdot \text{Weight}(f)/k)^k. \quad (2.6)$$

To prove it, let us use the inclusion-exclusion principle to compute $f(x)$. For each positive integer r , define $\psi_r: \{0, 1\}^n \rightarrow \mathbb{R}$ by

$$\psi_r(x) = \sum_{t=1}^r (-1)^{t-1} \sum_{\substack{S \subseteq [m] \\ |S|=t}} \bigwedge_{i \in S} f_i(x).$$

Since k is even, $\psi_k \leq f \leq \psi_{k-1}$, and we claim that

$$\mathbb{E}[\psi_{k-1} - \psi_k] \leq (e \cdot \text{Weight}(f)/k)^k. \quad (2.7)$$

Indeed, if $k > m$, then Equation (2.7) holds because $\psi_{k-1} \equiv \psi_k \equiv f$, and meanwhile if $k \leq m$, then

$$\begin{aligned} \mathbb{E}[\psi_{k-1} - \psi_k] &= \sum_{\substack{S \subseteq [m] \\ |S|=k}} \prod_{i \in S} \mathbb{E}[\phi_i] \leq \binom{m}{k} \cdot \left(\frac{\sum_{i=1}^m \mathbb{E}[\phi_i]}{m}\right)^k \leq \left(\frac{em}{k}\right)^k \cdot \left(\frac{\text{Weight}(f)}{m}\right)^k \\ &= (e \cdot \text{Weight}(f)/k)^k. \end{aligned}$$

(The first inequality follows from Maclaurin's inequality.) Thus, f is sandwiched between ψ_k and ψ_{k-1} . Furthermore, since the top gate of each ϕ_i is \wedge , each function $\bigwedge_{i \in S} \phi_i(x)$ is a depth- d read-once \mathbf{AC}^0 formula. Therefore, by the Triangle Inequality for PRG Errors, X fools ψ_r with error δ_r where

$$\delta_r = \varepsilon \cdot \sum_{t=1}^r \binom{m}{t} = \varepsilon \cdot \binom{m+r-1}{r} \leq \varepsilon \cdot (m+r)^r.$$

Since $\psi_r = \psi_m$ for all $r \geq m$, it follows that X fools ψ_r with error $\varepsilon \cdot (2m)^r$. Therefore, by the sandwiching lemma (Lemma 2.5.2), X fools ϕ with the error given by Equation (2.6).

Now let

$$k_* = \frac{\log(1/\varepsilon)}{2 \log(2m)},$$

or to be more precise, let k_* be the smallest even positive integer that is at least the above value. We split into two cases. For the first case, suppose $\text{Weight}(f) \leq k_*/(2e)$. Then we achieve error

$$\varepsilon \cdot (2m)^{k_*} + 2^{-k_*} = \exp\left(-\Omega\left(\frac{\log(1/\varepsilon)}{\log m}\right)\right).$$

Since f is read-once, $m \leq n$, so this error value is sufficient to establish the lemma. For the second case, suppose $\text{Weight}(f) > k_*/(2e)$. Let $f'(x) = f_1(x) \vee \cdots \vee f_{m'}(x)$, where m' is the largest value such that $\text{Weight}(f') \leq k_*/(2e)$. Then $f' \leq f \leq 1$, and

$$\mathbb{E}[1 - f'] = \prod_{i=1}^{m'} (1 - \mathbb{E}[f_i]) \leq e^{-\text{Weight}(f')} \leq e^{-\left(\frac{k_*}{2e} - 1\right)}.$$

Therefore, ϕ is δ -sandwiched between ϕ' and 1, where $\delta = \exp\left(-\Omega\left(\frac{\log(1/\varepsilon)}{\log m}\right)\right)$. Furthermore, X fools f' with error

$$\varepsilon \cdot (2m')^{k_*} + 2^{-k_*} \leq (2m)^{k_*} + 2^{-k_*} = \exp\left(-\Omega\left(\frac{\log(1/\varepsilon)}{\log m}\right)\right),$$

and obviously X fools 1 with error 0, so another application of the sandwiching lemma completes the proof in this case.

Finally, suppose the output gate of ϕ is \wedge . Then $\neg\phi$ can be computed by a depth- $(d+1)$ read-once formula where the output gate is \vee . Therefore, X fools $\neg\phi$, and hence it fools ϕ with the same error. \square

Remark 2.5.7. *More generally, we can consider any class \mathcal{F} of Boolean functions on n bits. (The interesting case is when \mathcal{F} is not closed under complement.) Let $\text{AND} \diamond \mathcal{F}$ denote the “read-once composition” of AND with \mathcal{F} , i.e., the class of functions of the form $f(x) = \bigwedge_{i=1}^t f_i(x)$ where $f_1, \dots, f_t \in \mathcal{F}$ and f_1, \dots, f_t depend on disjoint parts of the input. Define $\text{OR} \diamond \mathcal{F}$ similarly. The proof of Lemma 2.5.6 shows that if X fools $\text{AND} \diamond \mathcal{F}$ with error ε , then X fools $\text{OR} \diamond \mathcal{F}$ with error $\exp(-\Omega(\log(1/\varepsilon)/\log n))$.*

Proof of Theorem 2.5.5. By Proposition 2.3.7, if f is a depth-1 read-once \mathbf{AC}^0 formula, then either $L_1(f) \leq 1$ or $L_1(\neg f) \leq 1$. Either way, every δ -biased distribution fools f with error δ . This is the base case of an induction on d , where Lemma 2.5.6 is the inductive step. \square

We can also consider read- k depth- d \mathbf{AC}^0 circuits for $k > 1$. Servedio and Tan studied the case $d = 2$ [ST19b], improving on previous work by Klivans, Lee, and Wan [KLW10]. Both works show that small-bias distributions fool read- k CNFs and DNFs; in the case of polynomial-size DNFs, Servedio and Tan’s analysis [ST19b] leads to a seed length of $\text{poly}(k, \log(1/\varepsilon)) \cdot \log n$. The case of larger depth $d > 2$ is open.

Open Problem 2.5.8 (PRGs for read-twice \mathbf{AC}^0 circuits). *Design an explicit PRG for read-twice depth- d \mathbf{AC}^0 circuits with a better seed length than the state-of-the-art PRG for general depth- d \mathbf{AC}^0 circuits [Lyu22].*

2.5.4 The sandwiching lemma and the triangle inequality are always enough

Suppose we wish to show that every distribution that fools one class $\mathcal{F}_{\text{simp}}$ also fools another class \mathcal{F} . As discussed at the beginning of this section, we have presented two techniques for proving such a “transfer theorem”:

1. The first technique is to express each $f \in \mathcal{F}$ as a linear combination of functions in $\mathcal{F}_{\text{simp}}$ and invoke the Triangle Inequality for PRG Errors.
2. The second technique is to sandwich each $f \in \mathcal{F}$ between functions in $\mathcal{F}_{\text{simp}}$ and invoke the Sandwiching Lemma.

We will now show that *these are the only two techniques that are ever necessary*. That is, we will show that if every distribution that fools $\mathcal{F}_{\text{simp}}$ also fools \mathcal{F} , then that fact can be proven by sandwiching each $f \in \mathcal{F}$ between linear combinations of functions in $\mathcal{F}_{\text{simp}}$.

Theorem 2.5.9 (Characterization of when fooling one class implies fooling another). *Let $n \in \mathbb{N}$, let $\mathcal{F}_{\text{simp}}$ be a finite class of functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$, and let $g: \{0, 1\}^n \rightarrow \mathbb{R}$. Let $\varepsilon_0, \varepsilon > 0$ and suppose that every distribution X that fools $\mathcal{F}_{\text{simp}}$ with error ε_0 also fools g with error ε . Then g is (2ε) -sandwiched between two functions $f_\ell, f_u: \{0, 1\}^n \rightarrow \mathbb{R}$ of the form*

$$f_\ell(x) = \lambda_\ell^{(0)} + \sum_{i=1}^{k_\ell} \lambda_\ell^{(i)} f_\ell^{(i)}(x) \quad (2.8)$$

$$f_u(x) = \lambda_u^{(0)} + \sum_{i=1}^{k_u} \lambda_u^{(i)} f_u^{(i)}(x), \quad (2.9)$$

where $k_\ell, k_u \in \mathbb{N}$, $\lambda_\ell^{(i)}, \lambda_u^{(i)} \in \mathbb{R}$, $f_\ell^{(i)}, f_u^{(i)} \in \mathcal{F}_{\text{simp}}$, and

$$\varepsilon_0 \cdot \sum_{i=1}^{k_\ell} |\lambda_\ell^{(i)}| \leq \varepsilon \quad (2.10)$$

$$\varepsilon_0 \cdot \sum_{i=1}^{k_u} |\lambda_u^{(i)}| \leq \varepsilon. \quad (2.11)$$

Conversely, if we start from the assumption that [Equations \(2.8\) to \(2.11\)](#) hold, then for any distribution X that fools $\mathcal{F}_{\text{simp}}$ with error ε_0 , the Triangle Inequality for PRG Errors implies that X fools f_ℓ and f_u with error ε , and therefore the Sandwiching Lemma implies that X fools g with error 3ε . This recovers the assumption of [Theorem 2.5.9](#) up to a factor of three⁴ in the error parameter. In this sense, [Theorem 2.5.9](#) shows that the Triangle Inequality for PRG Errors and the Sandwiching Lemma are “complete.”

Before presenting the proof, let us elaborate on what the theorem says in two important special cases.

- Let $\mathcal{F}_{\text{simp}}$ be the class of Boolean k -juntas and let $\varepsilon_0 = 0$. Then [Theorem 2.5.9](#) says that a function is fooled by every k -wise uniform distribution if and only if the function can be sandwiched between two *low-degree polynomials*. This was first shown by Bazzi [[Baz09](#)].
- Next, let $\mathcal{F}_{\text{simp}}$ to be the class of parity functions. Then [Theorem 2.5.9](#) essentially says that a function is fooled by every small-bias distribution if and only if the function can be sandwiched between two functions with *low Fourier L_1 norm*.⁵ This was first shown by De, Etesami, Trevisan, and Tulsiani [[DETT10](#)].⁶

The general case seems to be folklore.

Proof of [Theorem 2.5.9](#). The proof uses linear programming duality. For each $f \in \mathcal{F}_{\text{simp}}$, define $\bar{f}: \{0, 1\}^n \rightarrow \mathbb{R}$ by $\bar{f}(x) = f(x) - \mathbb{E}[f]$. Consider the following linear program in the variables

⁴A more refined analysis, involving a more cumbersome version of the Sandwiching Lemma, gives a tight characterization without the extra factor of three.

⁵Actually the quantity that matters is the sum of absolute values of the *nonempty* Fourier coefficients, whereas we included the empty Fourier coefficient in our definition of Fourier L_1 norm.

⁶Note that there is a minor mistake in the formulation by De et al. [[DETT10](#)]: in their Proposition 2.7, the lower and upper sandwichers should be allowed to have different values of “ l ” and “ δ .”

$\{p_x\}_{x \in \{0,1\}^n}$:

$$\begin{aligned}
& \text{Maximize} && \sum_{x \in \{0,1\}^n} p_x g(x), \\
& \text{subject to} && p_x \geq 0 \text{ for all } x \in \{0,1\}^n \\
& && \text{and } \sum_{x \in \{0,1\}^n} p_x = 1 \\
& && \text{and } \sum_{x \in \{0,1\}^n} p_x \bar{f}(x) \leq \varepsilon_0 \text{ for all } f \in \mathcal{F}_{\text{simp}} \\
& && \text{and } - \sum_{x \in \{0,1\}^n} p_x \bar{f}(x) \leq \varepsilon_0 \text{ for all } f \in \mathcal{F}_{\text{simp}}.
\end{aligned}$$

The constraints say that the p_x variables are the probability mass function of some distribution that fools $\mathcal{F}_{\text{simp}}$ with error ε_0 . The program is feasible, because if nothing else we can set $p_x = 2^{-n}$ (the uniform distribution). The objective function is the expectation of g under the distribution defined by the p_x variables, so the optimal value must be at most $\mathbb{E}[g] + \varepsilon$.

The dual linear program, in the variables z and $\{y_f^+, y_f^-\}_{f \in \mathcal{F}_{\text{simp}}}$, is as follows:

$$\begin{aligned}
& \text{Minimize} && z + \varepsilon_0 \cdot \sum_{f \in \mathcal{F}_0} (y_f^+ + y_f^-), \\
& \text{subject to} && y_f^+, y_f^- \geq 0 \text{ for all } f \in \mathcal{F}_{\text{simp}} \\
& && \text{and } z + \sum_{f \in \mathcal{F}_0} \bar{f}(x) \cdot (y_f^+ - y_f^-) \geq g(x) \text{ for all } x \in \{0,1\}^n.
\end{aligned}$$

By strong LP duality, the optimal value of this dual linear program is also at most $\mathbb{E}[g] + \varepsilon$. Observe that given a feasible solution to the dual linear program, if we subtract $\min\{y_f^+, y_f^-\}$ from y_f^+ and from y_f^- , then we get another feasible solution and the objective function can only decrease. Therefore, by setting $y_f = y_f^+ - y_f^-$, we obtain real numbers z^* and $\{y_f^*\}_{f \in \mathcal{F}_{\text{simp}}}$ such that

$$\begin{aligned}
& z^* + \varepsilon_0 \cdot \sum_{f \in \mathcal{F}_{\text{simp}}} |y_f^*| \leq \mathbb{E}[g] + \varepsilon, \text{ and} \\
& z^* + \sum_{f \in \mathcal{F}_{\text{simp}}} \bar{f}(x) y_f^* \geq g(x) \text{ for all } x \in \{0,1\}^n.
\end{aligned}$$

Define

$$\begin{aligned}
f_u(x) &= z^* + \sum_{f \in \mathcal{F}_{\text{simp}}} y_f^* \cdot \bar{f}(x) \\
&= \left(z^* - \sum_{f \in \mathcal{F}_{\text{simp}}} y_f^* \mathbb{E}[f] \right) + \sum_{f \in \mathcal{F}_{\text{simp}}} y_f^* \cdot f(x).
\end{aligned}$$

Then f_u has the form given by [Equation \(2.9\)](#), and $f_u \geq g$. Furthermore, $\mathbb{E}[f_u] = z^*$, so

$$0 \leq \mathbb{E}[f_u - g] = z^* - \mathbb{E}[g] \leq \varepsilon - \varepsilon_0 \cdot \sum_{f \in \mathcal{F}_{\text{simp}}} |y_f^*|.$$




Figure 2.5: A depth-4 size-13 \mathbf{AC}^0 circuit computing the function $f(x) = \text{MAJ}(x_1 \oplus x_2, x_3 \oplus x_4, x_5 \oplus x_6)$.

This shows that $\mathbb{E}[f_u - g] \leq \varepsilon$ and that Equation (2.11) holds.

Fooling g is equivalent to fooling $-g$, so the above also shows that there is some function f_ℓ of the form given by Equation (2.8) such that $-f_\ell \geq -g$, $\mathbb{E}[g - f_\ell] \leq \varepsilon$, and Equation (2.10) holds. Therefore, g is (2ε) -sandwiched between f_ℓ and f_u . \square

2.6 Braverman’s theorem: Limited independence fools \mathbf{AC}^0

In Section 2.5.3, we presented a PRG for read-once \mathbf{AC}^0 formulas. In this section, we will present a PRG for general (read-many) \mathbf{AC}^0 circuits (see Figure 2.5). In particular, we will show that every k -wise uniform generator fools constant-depth polynomial-size \mathbf{AC}^0 circuits for a suitable $k = \text{polylog}(n)$. This was first conjectured by Linial and Nisan [LN90]. Two decades later, it was proved to be true for depth-2 circuits by Bazzi [Baz09] and a simpler proof of this was discovered by Razborov [Raz09]. Building on this line of work, Braverman [Bra10] proved that k -wise independence for polylogarithmic k fools \mathbf{AC}^0 circuits. The parameters were subsequently improved by Tal [Tal17] and Harsha and Srinivasan [HS19], leading to the following.

Theorem 2.6.1 (Braverman’s theorem [Bra10; Tal17; HS19]). *For every $n, m, d \in \mathbb{N}$ and $\varepsilon > 0$, there is a value $k = (\log m)^{O(d)} \cdot \log(1/\varepsilon)$ such that if X is a k -wise uniform distribution over $\{0, 1\}^n$, then X fools size- m depth- d \mathbf{AC}^0 circuits with error ε . Consequently, there is an explicit ε -PRG for depth- d size- m \mathbf{AC}^0 circuits with seed length*

$$(\log m)^{O(d)} \cdot \log n \cdot \log(1/\varepsilon).$$

For context, Braverman’s theorem represents neither the first nor the best unconditional PRG for \mathbf{AC}^0 known. Instead, the advantage of Braverman’s theorem is that k -wise uniformity is a particularly *simple and general* PRG construction. That being said, the *analysis* is quite nontrivial, as we will see.

We will present two proofs of Braverman’s theorem. First, we present a novel proof that is arguably simpler⁷ than previous proofs, but it does not give the best parameters – the value of

⁷In particular, the new proof does not require “probabilistic polynomials” for \mathbf{AC}^0 circuits.

k will be slightly worse than what [Theorem 2.6.1](#) promises. Then, we will present the known state-of-the-art proof [[Bra10](#); [Tal17](#); [HS19](#)].

2.6.1 LMN polynomials

Observe that every degree- k polynomial over the reals is perfectly fooled by k -wise uniform generators. To prove that k -wise uniform generators fool \mathbf{AC}^0 circuits, our approach will be to show that \mathbf{AC}^0 circuits are sandwiched between degree- k polynomials. Our starting point is the Linial-Mansour-Nisan theorem [[LMN93](#)] and its subsequent improvements [[Bop97](#); [Hås01](#); [Tal17](#)], which show that \mathbf{AC}^0 circuits can indeed be *approximated* by bounded low-degree polynomials in the L_2 norm.

Theorem 2.6.2 (LMN polynomials [[Tal17](#)]). *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be computable by a size- m depth- d \mathbf{AC}^0 circuit and let $\gamma > 0$. There exists $\tilde{f}: \{0, 1\}^n \rightarrow \mathbb{R}$ such that:*

1. (*L_2 approximation*) We have

$$\|f - \tilde{f}\|_2^2 \stackrel{\text{def}}{=} \mathbb{E}_{x \sim U_n} \left[|f(x) - \tilde{f}(x)|^2 \right] \leq \gamma. \quad (2.12)$$

2. (*Low-degree*) We have

$$\deg(\tilde{f}) \leq O(\log m)^{d-1} \cdot \log(1/\gamma).$$

3. (*Bounded*) For every $x \in \{0, 1\}^n$, we have

$$|\tilde{f}(x)| \leq 2^{O(\log m)^{d-1} \cdot \log(1/\gamma) \cdot \log \log m}. \quad (2.13)$$

The proof of [Theorem 2.6.2](#) uses *random restrictions and switching lemmas* to analyze the *Fourier spectrum* of \mathbf{AC}^0 circuits.⁸ We will not study the proof of [Theorem 2.6.2](#) here. Instead, we will take [Theorem 2.6.2](#) for granted, and use it to show that \mathbf{AC}^0 circuits are sandwiched by low-degree polynomials.

2.6.2 Operations on functions with low-degree sandwiching polynomials

We begin with a couple of lemmas about sandwiching by low-degree polynomials.

Lemma 2.6.3 (If f and g have low-degree sandwichers, then $f + g$ has low-degree sandwichers). *Suppose that $f: \{0, 1\}^n \rightarrow \mathbb{R}$ is ε -sandwiched by polynomials of degree k and $g: \{0, 1\}^n \rightarrow \mathbb{R}$ is δ -sandwiched by polynomials of degree k . Then the sum $f + g$ is $(\varepsilon + \delta)$ -sandwiched by polynomials of degree k .*

We omit the simple proof.

Lemma 2.6.4 (If f has low-degree sandwichers and g is a bounded low-degree polynomial, then $f \cdot g$ has low-degree sandwichers). *Suppose that $f: \{0, 1\}^n \rightarrow \mathbb{R}$ is ε -sandwiched by polynomials of degree k . Let $g: \{0, 1\}^n \rightarrow [-L, L]$, and let $h(x) = f(x) \cdot g(x)$. Then h is $O(\varepsilon \cdot L)$ -sandwiched by polynomials of degree $k + \deg(g)$.*

⁸The polynomial \tilde{f} is defined by dropping all but the lowest-degree Fourier coefficients of f . The bounds [Equations \(2.12\)](#) and [\(2.13\)](#) follow from a bound on the L_2 tail of f [[Tal17](#), Theorem 1] and a bound on the L_1 growth of f [[Tal17](#), Theorem 37].

Proof. Let f_ℓ, f_u be the sandwichers for f . Suppose first that g is $[0, L]$ -valued. In this case, for every $x \in \{0, 1\}^n$,

$$f_\ell(x) \cdot g(x) \leq h(x) \leq f_u(x) \cdot g(x),$$

and

$$(f_u(x) - f_\ell(x)) \cdot g(x) \leq L \cdot (f_u(x) - f_\ell(x)),$$

showing that h is $(\varepsilon \cdot L)$ -sandwiched between $f_\ell \cdot g$ and $f_u \cdot g$, each of which has degree $k + \deg(g)$.

Next, suppose that g is $[-L, 0]$ -valued. The previous argument shows that $-fg$ is $(\varepsilon \cdot L)$ -sandwiched by polynomials of degree $k + \deg(g)$, and therefore so is fg by negating *and swapping* the sandwichers.

Finally, consider the general case that g is $[-L, L]$ -valued. Write $g = -L + g'$, where g' is $[0, 2L]$ -valued. Then

$$h = -L \cdot f + f \cdot g'.$$

By our previous analyses, $-L \cdot f$ is $(\varepsilon \cdot L)$ -sandwiched by polynomials of degree k , and $f \cdot g'$ is $(2\varepsilon \cdot L)$ -sandwiched by polynomials of degree $k + \deg(g') = k + \deg(g)$. By [Lemma 2.6.3](#), it follows that h is $(3\varepsilon \cdot L)$ -sandwiched by polynomials of degree $k + \deg(g)$. \square

2.6.3 Low-degree sandwichers for \mathbf{AC}^0 circuits

We are now prepared to show that \mathbf{AC}^0 circuits are sandwiched by low-degree polynomials, and hence they are fooled by k -wise uniform distributions.

Theorem 2.6.5 (\mathbf{AC}^0 circuits are sandwiched by low-degree polynomials). *Let $m, d \in \mathbb{N}$ and $\varepsilon > 0$. Every size- m depth- d \mathbf{AC}^0 circuit f is ε -sandwiched by polynomials of degree $(\log m)^{O(d^2)} \cdot \log(1/\varepsilon)$.*

By the Sandwiching Lemma, [Theorem 2.6.5](#) implies Braverman's Theorem ([Theorem 2.6.1](#)), albeit with $k = (\log m)^{O(d^2)} \cdot \log(1/\varepsilon)$ instead of $k = (\log m)^{O(d)} \cdot \log(1/\varepsilon)$.

Proof of [Theorem 2.6.5](#). We will show by induction on d that f has ε -sandwiching polynomials of degree

$$(c \log m)^{(d^2-d)/2} \cdot (\log \log m)^{d-1} \cdot \lceil \log(m/\varepsilon) \rceil$$

for a suitable constant c . First, consider the base case $d = 1$, and assume without loss of generality that f is an AND of m literals.

- If $m < \log(1/\varepsilon)$, then f can be computed exactly by a polynomial of degree m , so we are done.
- If $m \geq \log(1/\varepsilon)$, then our upper sandwicher is the product of the first $\lceil \log(1/\varepsilon) \rceil$ literals and our lower sandwicher is the constant 0 function.

Now, for the inductive step, suppose f has depth $d > 1$. Assume without loss of generality that the top gate of f is OR, say $f = \bigvee_{i=1}^m f_i$. We reason similarly to Bazzi [[Baz09](#)] and Razborov [[Raz09](#)]. For each $i \in [m]$, let $g_i = \bigwedge_{j=1}^{i-1} (\neg f_j)$, so that

$$f = \sum_{i=1}^m f_i \cdot g_i.$$

Each function g_i is a size- m depth- d \mathbf{AC}^0 circuit; let \tilde{g}_i be the corresponding polynomial approximation from [Theorem 2.6.2](#) with error parameter $\gamma = \varepsilon/(2m^3)$. We define

$$h = \sum_{i=1}^m f_i \cdot \tilde{g}_i$$

$$f_\ell = f - (f - h)^2 \tag{2.14}$$

$$f_u = f + (f - h)^2 \cdot \left(\left(\sum_{i=1}^m f_i \right) - f \right). \tag{2.15}$$

The plan is, we will show that f is sandwiched between f_ℓ and f_u , and then we will use our induction hypothesis to show that f_ℓ and f_u are sandwiched by low-degree polynomials. Consequently, f itself is sandwiched by low-degree polynomials.

From the definitions, clearly $f_\ell \leq f \leq f_u$. Furthermore,

$$\begin{aligned} \mathbb{E}[f_u - f_\ell] &= \mathbb{E} \left[(f - h)^2 \cdot \left(1 - f + \sum_{i=1}^m f_i \right) \right] \leq m \cdot \mathbb{E}[(f - h)^2] \\ &= m \cdot \mathbb{E} \left[\left(\sum_{i=1}^m f_i \cdot (g_i - \tilde{g}_i) \right)^2 \right] \\ &\leq m^2 \cdot \sum_{i=1}^m \mathbb{E}[f_i^2 \cdot (g_i - \tilde{g}_i)^2] \\ &\leq m^3 \cdot \gamma = \varepsilon/2. \end{aligned}$$

Now we turn to showing that f_ℓ and f_u are themselves sandwiched by low-degree polynomials. For the first step, we claim that

$$f_\ell = 1 - (1 - h)^2 \tag{2.16}$$

$$f_u = 1 + (1 - h)^2 \cdot \left(\left(\sum_{i=1}^m f_i \right) - 1 \right). \tag{2.17}$$

Indeed, when $f(x) = 1$, this is clear, since we have simply substituted 1 for each appearance of f in [Equations \(2.14\)](#) and [\(2.15\)](#). Meanwhile, when $f(x) = 0$, [Equations \(2.16\)](#) and [\(2.17\)](#) still hold, because $f_i(x) = 0$ for every i and therefore $\sum_{i=1}^m f_i(x) = 0$ and $h(x) = 0$.

Next, we will plug the definition of h into [Equations \(2.16\)](#) and [\(2.17\)](#) and expand. For convenience, define $f_0 = \tilde{g}_0 = 1$. That way, we get

$$f_\ell = \sum_{i=0}^m \sum_{j=0}^m c_{i,j} \cdot f_i \cdot f_j \cdot \tilde{g}_i \cdot \tilde{g}_j$$

$$f_u = \sum_{i=0}^m \sum_{j=0}^m \sum_{k=0}^m c_{i,j,k} \cdot f_i \cdot f_j \cdot f_k \cdot \tilde{g}_j \cdot \tilde{g}_k,$$

where $|c_{i,j}| \leq 1$ and $|c_{i,j,k}| \leq 1$. For simplicity, let us focus on a single term from the expansion of f_u , namely a term of the form $c_{i,j,k} \cdot f_i \cdot f_j \cdot f_k \cdot \tilde{g}_j \cdot \tilde{g}_k$. By [Equation \(2.13\)](#), for every x , we have $|c_{i,j,k} \cdot \tilde{g}_j(x) \cdot \tilde{g}_k(x)| \leq L$ where

$$L = 2^{O(\log m)^{d-1} \cdot \log(m/\varepsilon) \cdot \log \log m}.$$

Now, each subcircuit f_i has an AND gate on top, so the product $f_i \cdot f_j \cdot f_k$ can be computed by a size- m depth- $(d-1)$ \mathbf{AC}^0 circuit. Therefore, by induction, for every $\delta > 0$, the product $f_i \cdot f_j \cdot f_k$ is δ -sandwiched by polynomials of degree D where

$$D = (c \log m)^{((d-1)^2 - (d-1))/2} \cdot (\log \log m)^{d-2} \cdot \lceil \log(m/\delta) \rceil.$$

We select $\delta = \Theta(\frac{\varepsilon}{Lm^3})$. That way, [Lemma 2.6.4](#) ensures that the term $c_{i,j,k} \cdot f_i \cdot f_j \cdot f_k \cdot \tilde{g}_j \cdot \tilde{g}_k$ is $(\frac{\varepsilon}{4(m+1)^3})$ -sandwiched by polynomials of degree $D + \deg(\tilde{g}_j \cdot \tilde{g}_k)$. Therefore, by [Lemma 2.6.3](#), f_u as a whole (and similarly f_ℓ as well) is $(\varepsilon/4)$ -sandwiched by polynomials of degree $D + \deg(\tilde{g}_j \cdot \tilde{g}_k)$. Consequently, f is ε -sandwiched by polynomials of degree $D + \deg(\tilde{g}_j \cdot \tilde{g}_k)$. All that remains is to simplify the degree bound:

$$\begin{aligned} D + \deg(\tilde{g}_j \cdot \tilde{g}_k) &= (c \log m)^{((d-1)^2 - (d-1))/2} \cdot (\log \log m)^{d-2} \cdot \lceil \log(m/\delta) \rceil + O(\log m)^{d-1} \cdot \log(1/\gamma) \\ &= (c \log m)^{((d-1)^2 - (d-1))/2} \cdot O(\log m)^{d-1} \cdot (\log \log m)^{d-1} \cdot \log(m/\varepsilon) \\ &\leq (c \log m)^{(d^2 - d)/2} \cdot (\log \log m)^{d-1} \cdot \lceil \log(m/\varepsilon) \rceil, \end{aligned}$$

provided that we choose the constant c large enough. (Note that $\frac{(d-1)^2 - (d-1)}{2} + d - 1 = \frac{d^2 - d}{2}$.) \square

2.6.4 Improved parameters via probabilistic polynomials

So far, we have shown that k -wise uniform generators fool size- m depth- d \mathbf{AC}^0 circuits where $k = (\log m)^{O(d^2)} \cdot \log(1/\varepsilon)$. Next, we will show how to improve the exponent from $O(d^2)$ to $O(d)$.⁹ The improvement relies on a line of work constructing *probabilistic real polynomials* for \mathbf{AC}^0 circuits [[BRS91](#); [Tar93](#); [Bra10](#); [HS19](#)], starting with two independent papers by Beigel, Reingold, and Spielman [[BRS91](#)] and Tarui [[Tar93](#)] (see also, e.g., work by Razborov [[Raz87](#)], Smolensky [[Smo87](#)], and Toda and Ogiwara [[TO92](#)]). These works show that for every \mathbf{AC}^0 circuit f , there is a distribution F over low-degree polynomials such that for each fixed input $x \in \{0, 1\}^n$, with high probability over $\tilde{f} \sim F$, we have the exact equality $\tilde{f}(x) = f(x)$. In our setting, we will actually be thinking of the input x as random, which allows us to fix some $\tilde{f} \in \text{Supp}(F)$ that agrees with f with high probability over the choice of input. Furthermore, even in the “bad case” that $f(x) \neq \tilde{f}(x)$, the constructions still have some guarantees. The best parameters known are achieved by Harsha and Srinivasan [[HS19](#)], who prove the following:

Theorem 2.6.6 (BRS-Tarui polynomials [[HS19](#)]). *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be computable by a depth- d size- m \mathbf{AC}^0 circuit, let $\delta > 0$, and let D be a distribution over $\{0, 1\}^n$. There exist a polynomial $\tilde{f}: \{0, 1\}^n \rightarrow \mathbb{R}$ and an “error function” $E: \{0, 1\}^n \rightarrow \{0, 1\}$ such that*

- $\mathbb{E}[E(D)] \leq \delta$, and if $E(x) = 0$, then $f(x) = \tilde{f}(x)$. (Hence, $\Pr_{x \sim D} [f(x) \neq \tilde{f}(x)] \leq \delta$.)
- $\deg(\tilde{f}) \leq (\log m)^{O(d)} \log(1/\delta)$ and $\|\tilde{f}\|_\infty \stackrel{\text{def}}{=} \max_x |\tilde{f}(x)| \leq \exp((\log m)^{O(d)} \log(1/\delta))$.
- E can be computed by an \mathbf{AC}^0 circuit of size¹⁰ $m^{O(1)}$ and depth $d + O(1)$.

⁹Note that for certain small values of d such as $d = 3$, the parameters from the first proof are actually superior to the parameters from the second proof. We thank Avishay Tal for pointing this out (personal communication).

¹⁰In Harsha and Srinivasan’s work [[HS19](#)], the size bound is stated as $(m \log(1/\delta))^{O(1)}$. We may assume without loss of generality that $\log(1/\delta) < m$, because f can be computed exactly by a degree- m polynomial (namely its Fourier expansion).

Note that [Theorem 2.6.6](#) provides a low-degree approximation over an arbitrary input distribution, unlike LMN polynomials ([Theorem 2.6.2](#)) which are specific to the uniform distribution. The constructions of probabilistic polynomials for \mathbf{AC}^0 [[BRS91](#); [Tar93](#); [Bra10](#); [HS19](#)] rely on Valiant and Vazirani's *isolation lemma* [[VV86](#)].

Braverman's original proof that k -wise uniformity fools \mathbf{AC}^0 circuits for polylogarithmic k was based on a clever combination of LMN polynomials and BRS-Tarui polynomials. We present (a version of) that proof below to prove [Theorem 2.6.1](#).

Lemma 2.6.7 (Sandwichers for \mathbf{AC}^0 with better parameters). *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be computable by a depth- d size- m \mathbf{AC}^0 circuit, let $\delta > 0$, and let D be a distribution over $\{0, 1\}^n$. There exist polynomials $p_\ell, p_u: \{0, 1\}^n \rightarrow \mathbb{R}$ of degree $(\log m)^{O(d)} \log(1/\delta)$ and a function $E: \{0, 1\}^n \rightarrow \{0, 1\}$ such that f is δ -sandwiched between $p_\ell - E$ and $p_u + E$, and furthermore,*

$$\mathbb{E}[E(U_n)] \leq \delta \text{ and } \mathbb{E}[E(D)] \leq \delta. \quad (2.18)$$

To be clear, f is sandwiched between $p_\ell - E$ and $p_u + E$ with respect to the uniform distribution (see [Definition 2.5.1](#)). The only part of the conclusion that says something about the arbitrary distribution D is the bound $\mathbb{E}[E(D)] \leq \delta$.

Proof. Let $D' = \frac{1}{2}(D + U_n)$, i.e., D' is D with probability $1/2$ and U_n with probability $1/2$. Let \tilde{f} be the BRS-Tarui polynomial for f from [Theorem 2.6.6](#) and let E be the corresponding error function with respect to the distribution D' with error $\mathbb{E}[E(D')] \leq \delta/24$. Note that

$$\mathbb{E}[E(D')] = \frac{1}{2} \mathbb{E}[E(D)] + \frac{1}{2} \mathbb{E}[E(U_n)] \leq \frac{\delta}{24},$$

so $\mathbb{E}[E(D)] \leq \delta/12$ and $\mathbb{E}[E(U_n)] \leq \delta/12$, proving [Equation \(2.18\)](#).

Recall that the error function E can be computed by an \mathbf{AC}^0 circuit of size $m^{O(1)}$ and depth $d + O(1)$ (see [Theorem 2.6.6](#)). Therefore, we may apply [Theorem 2.6.2](#) to get an LMN polynomial \tilde{E} that satisfies $\|E - \tilde{E}\|_2^2 \leq \gamma$ for an error parameter γ that will be specified later. Define three more approximations to f by the formulas

$$\begin{aligned} \phi &= 1 - (1 - f) \cdot (1 - E) = f \vee E & (2.19) \\ \tilde{\phi} &= 1 - (1 - \tilde{f}) \cdot (1 - \tilde{E}) \\ p_u &= \left(1 - (1 - \tilde{f}) \cdot (1 - \tilde{E})\right)^2 = \left(\tilde{\phi}\right)^2. \end{aligned}$$

We must show that $p_u + E$ is an upper sandwicher for f . By a case analysis, let us prove that the following two bounds hold (pointwise):

$$f \leq p_u + E \quad (2.20)$$

$$p_u \leq f + 2E + 2 \cdot \left(\tilde{\phi} - \phi\right)^2. \quad (2.21)$$

- (Case 1) Suppose $f(x) > E(x)$, i.e., $E(x) = 0$ and $f(x) = 1$. Then $\tilde{f}(x) = f(x) = 1$, so $p_u(x) = 1$ as well, which implies [Equations \(2.20\)](#) and [\(2.21\)](#).
- (Case 2) Suppose $f(x) \leq E(x)$, i.e., $E(x) = 1$ or $f(x) = 0$. Then [Equation \(2.20\)](#) holds because p_u is non-negative. Furthermore, $\phi(x) = E(x)$ in this case, so

$$p_u(x) = \left(E(x) + \tilde{\phi}(x) - \phi(x)\right)^2 \leq 2E(x)^2 + 2 \cdot \left(\tilde{\phi}(x) - \phi(x)\right)^2,$$

which implies [Equation \(2.21\)](#) because $E(x)^2 = E(x)$.

Now, because of the factor of $1 - E$ in Equation (2.19), we have the identity $\phi = 1 - (1 - \tilde{f}) \cdot (1 - E)$, so Equation (2.21) becomes

$$p_u \leq f + 2E + 2 \cdot \left((1 - \tilde{f}) \cdot (E - \tilde{E}) \right)^2.$$

Therefore,

$$\begin{aligned} \mathbb{E}[p_u + E - f] &\leq 3\mathbb{E}[E] + 2\mathbb{E} \left[\left((1 - \tilde{f}) \cdot (E - \tilde{E}) \right)^2 \right] \\ &\leq \delta/4 + 2 \cdot (1 + \|\tilde{f}\|_\infty)^2 \cdot \|E - \tilde{E}\|_2^2 \\ &\leq \delta/4 + 2 \cdot (1 + \|\tilde{f}\|_\infty)^2 \cdot \gamma. \end{aligned}$$

By choosing $\gamma = \delta / (8 \cdot (1 + \|\tilde{f}\|_\infty)^2)$, we get $\mathbb{E}[p_u + E - f] \leq \delta/2$.

Next, let us bound the degree of p_u . Recall that Theorems 2.6.2 and 2.6.6 give the bounds

$$\begin{aligned} \deg(\tilde{f}) &\leq (\log m)^{O(d)} \log(1/\delta) \\ \|\tilde{f}\|_\infty &\leq \exp \left((\log m)^{O(d)} \log(1/\delta) \right) \\ \deg(\tilde{E}) &\leq O(\log m)^{d+O(1)} \log(1/\gamma) = (\log m)^{O(d)} \log(1/\delta). \end{aligned}$$

Therefore, $\deg(p_u) \leq 2\deg(\tilde{f}) + 2\deg(\tilde{E}) \leq (\log m)^{O(d)} \log(1/\delta)$.

To summarize, we have shown that every size- m depth- d \mathbf{AC}^0 circuit f can be upper-sandwiched by $p_u + E$ where p_u is a low-degree polynomial and E is a Boolean function with low expectation under both U_n and D . The class of size- m depth- d \mathbf{AC}^0 circuits is closed under complementation, so $1 - f$ has an upper sandwicher of the same form. Therefore, f can be lower-sandwiched by $p_\ell - E'$ where p_ℓ is a low-degree polynomial and E' is a Boolean function with low expectation under both U_n and D . If we use $1 - \tilde{f}$ as our BRS-Tarui polynomial for $1 - f$ in the above argument, then we can furthermore ensure $E' \equiv E$. \square

Braverman's theorem follows readily from Lemma 2.6.7:

Proof of Theorem 2.6.1. Let D be a k -wise independent distribution, where k is the bound on the degrees of the low-degree polynomials p_u and p_ℓ from Lemma 2.6.7 with $\delta = \varepsilon/2$. By Lemma 2.6.7, the circuit f is $(\varepsilon/2)$ -sandwiched between $p_\ell - E$ and $p_u + E$, where E is $(\varepsilon/2)$ -fooled by D . Since p_u and p_ℓ are degree- k polynomials, they are perfectly fooled by D . Therefore, by the sandwiching lemma, f is ε -fooled by D . \square

Intriguingly, although the proof of Theorem 2.6.1 is a sandwiching argument, the sandwichers are apparently *not* low-degree polynomials. They are of the form $p_\ell - E$ and $p_u + E$, where p_ℓ and p_u are low-degree polynomials, but the ‘‘error function’’ E does not seem to be a low-degree polynomial. (Furthermore, the sandwichers depend on the pseudorandom distribution D .)

In a formal sense, sandwiching polynomials are the only tool one ever needs to prove that k -wise uniform generators fool some class of functions, as discussed in Section 2.5.4. However, the proof of Braverman's theorem demonstrates that in practice, it is wise to ‘‘think outside the box’’ and consider other, more creative arguments. The technique of designing low-complexity error indicator functions, along the lines of Theorem 2.6.6, has turned out to be useful in other PRG problems [MRT19; DHH20; HHTT22].

It is an open problem to improve the parameters of Braverman's theorem even further. What is the optimal k such that every k -wise uniform generator fools depth- d size- m \mathbf{AC}^0 circuits with

error ε ? There are counterexamples showing that $k = \Omega((\log m)^{d-1} \log(1/\varepsilon))$ [LV96], but that still leaves a significant gap between the lower and upper bounds.

Open Problem 2.6.8 (Improved parameters for Braverman's theorem). *Show that for every $m, d \in \mathbb{N}$ and $\varepsilon > 0$, there exists a value*

$$k = (\log m)^{d+O(1)} \log(1/\varepsilon)$$

such that every k -wise uniform distribution fools depth- d size- m \mathbf{AC}^0 circuits with error ε .

Explicit PRGs for \mathbf{AC}^0 circuits with seed length $(\log m)^{d+O(1)} \log(1/\varepsilon)$ are already known (see [Section 5.3](#)); the question is whether a generic k -wise uniform generator does the job.

Chapter 3

Recycling Random Bits

In this chapter, we will present a few PRG constructions based on the paradigm of *recycling* random bits. In its simplest form, the idea is that we start by sampling $n/2$ *truly* random bits X . Then we “mix in” a few more truly random bits in some way, producing $n/2$ additional bits Y . Our final output is the concatenation (X, Y) .

This type of approach tends to make the most sense if there is a “communication bottleneck” between the part of the computation that processes X and the part of the computation that processes Y . Indeed, our first instantiation of this paradigm will be a PRG for two-party communication protocols in the next section.

3.1 PRGs for two-party communication protocols

In this section, we will present a PRG that fools two-party interactive communication protocols.

Definition 3.1.1 (Two-party communication protocol). *Let n be an even positive integer. In a two-party protocol on n bits with communication cost m , Alice holds $x \in \{0, 1\}^{n/2}$ and Bob holds $y \in \{0, 1\}^{n/2}$. They communicate interactively; in each round, the identity of the speaker is a function of all the bits transmitted so far, and the content of the message is a function of the bits transmitted and that party’s input (x or y). After at most m bits have been transmitted in total, the protocol terminates, and both parties output the same bit $f(x, y)$.*

Are these protocols deterministic, or are they randomized? Both, in a sense: the protocol is deterministic after fixing x and y , but we are thinking of x and y as random bits and seeking to replace them with pseudorandom bits. The seed length in the following theorem, due to Impagliazzo, Nisan, and Wigderson [INW94], is optimal.

Theorem 3.1.2 ([INW94]). *For every $n, m \in \mathbb{N}$ and every $\varepsilon > 0$, there is an explicit ε -PRG for two-party protocols on n bits with communication cost m with seed length $\frac{n}{2} + O(m + \log(1/\varepsilon))$.*

3.1.1 Expander graphs from a PRG perspective

Toward proving Theorem 3.1.2, let us take a detour to briefly discuss *expander graphs*. Doing justice to the topic of expanders is beyond the scope of this work; see, e.g., Vadhan’s monograph [Vad12] for a thorough treatment. Let us recall the definition.

Definition 3.1.3 (Spectral expanders). *Let G be a regular undirected graph with transition probability matrix $M \in [0, 1]^{N \times N}$. Let the eigenvalues of M be $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N$. We say G is an ε -spectral expander if $|\lambda_i| \leq \varepsilon$ for $i = 2, 3, \dots, N$.*

There is an alternative and equivalent definition of expander graphs in the language of PRGs. We focus on the special case that the number of vertices is a power of two, because that's the most natural scenario from the PRG perspective, but the following lemma generalizes in the natural way to an arbitrary number of vertices.

Lemma 3.1.4 (PRG characterization of spectral expander graphs). *Let n be an even positive integer, and let \mathcal{F} be the class of functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$ of the form $f(x, y) = g(x) \cdot h(y)$ where $g, h: \{0, 1\}^{n/2} \rightarrow \mathbb{R}$ satisfy $\text{Var}[g] \leq 1$ and $\text{Var}[h] \leq 1$. Let G be a regular undirected graph on the vertex set $\{0, 1\}^{n/2}$, sample a uniform random vertex X in G , and sample a uniform random neighbor Y of X . For every $\varepsilon > 0$, the following are equivalent.*

1. G is an ε -spectral expander (see [Definition 3.1.3](#)).
2. The distribution (X, Y) fools \mathcal{F} with error ε .

Proof. We will identify functions mapping $\{0, 1\}^{n/2} \rightarrow \mathbb{R}$ with column vectors in the space $\mathbb{R}^{2^{n/2}}$ endowed with the inner product $\langle g, h \rangle = \mathbb{E}_{U \sim U_{n/2}}[g(U) \cdot h(U)]$ and the norm $\|g\| = \sqrt{\langle g, g \rangle}$.

(1 \implies 2) Fix an arbitrary $f \in \mathcal{F}$, say $f(x, y) = g(x) \cdot h(y)$ where $\text{Var}[g] \leq 1$ and $\text{Var}[h] \leq 1$. First, suppose $\mathbb{E}[g] = \mathbb{E}[h] = 0$. Then

$$|\mathbb{E}[f(X, Y)]| = |\langle g, Mh \rangle| \leq \|g\| \cdot \|Mh\|$$

by the Cauchy-Schwarz inequality. Since G is regular, the all-ones vector is an eigenvector of M with eigenvalue $\lambda_1 = 1$. Since $\mathbb{E}[g] = 0$, the vector g is orthogonal to the all-ones vector. By the spectral theorem for real symmetric matrices, it follows that g is a linear combination of eigenvectors *other* than the all-ones vector, so by [Definition 3.1.3](#), $\|Mg\| \leq \varepsilon \cdot \|g\|$. Therefore,

$$|\mathbb{E}[f(X, Y)]| \leq \varepsilon \cdot \|g\| \cdot \|h\| = \varepsilon \cdot \sqrt{\text{Var}[g]} \cdot \sqrt{\text{Var}[h]} \leq \varepsilon.$$

Now for the general case, write $g = \mathbb{E}[g] + \bar{g}$ and $h = \mathbb{E}[h] + \bar{h}$, so

$$f(x, y) = \mathbb{E}[g] \cdot \mathbb{E}[h] + \bar{g}(x) \cdot \bar{h}(y) + \mathbb{E}[g] \cdot \bar{h}(y) + \mathbb{E}[h] \cdot \bar{g}(x).$$

Since G is regular, each marginal distribution X and Y is uniform over $\{0, 1\}^{n/2}$. Therefore, $\mathbb{E}[\bar{g}(X)] = \mathbb{E}[\bar{h}(Y)] = 0$. Furthermore, $\text{Var}[\bar{g}] = \text{Var}[g] \leq 1$ and $\text{Var}[\bar{h}] = \text{Var}[h] \leq 1$, so

$$|\mathbb{E}[f(X, Y)] - \mathbb{E}[f]| = |\mathbb{E}[\bar{g}(X) \cdot \bar{h}(Y)]| \leq \varepsilon.$$

(2 \implies 1) Let g be a unit eigenvector of M with eigenvalue λ and assume that g is not parallel to the all-ones vector. By the spectral theorem for symmetric matrices, it follows that g is orthogonal to the all-ones vector, i.e., $\mathbb{E}[g] = 0$. Since g has unit norm as a vector, we have $\text{Var}[g] = 1$. Therefore,

$$|\lambda| = |\langle Mg, g \rangle| = |\mathbb{E}[g(X) \cdot g(Y)]| \leq \varepsilon. \quad \square$$

The seed length required to sample the distribution (X, Y) in [Lemma 3.1.4](#) is governed by the *degree* of the expander graph. There exist explicit expanders with degree $\text{poly}(1/\varepsilon)$:

Theorem 3.1.5 (Explicit expanders). *For every $n \in \mathbb{N}$ and every $\varepsilon > 0$, there exists an ε -spectral expander with vertex set $\{0, 1\}^n$ and degree $D = \text{poly}(1/\varepsilon)$. The expander is “strongly explicit,” i.e., given a vertex x and a value $i \in [D]$, the i^{th} neighbor of x can be computed in time $\text{poly}(n, \log(1/\varepsilon))$.*

See Vadhan’s monograph [\[Vad12\]](#) for a discussion of approaches to proving [Theorem 3.1.5](#). [Theorem 3.1.5](#) translates to a seed length of $n/2 + O(\log(1/\varepsilon))$ in [Lemma 3.1.4](#). Pushing to the extreme, a “Ramanujan graph” is an ε -spectral expander of degree D where $\varepsilon = 2\sqrt{D} - 1$, which is essentially the best possible [\[Alo86; Nil91; Fri93\]](#) and which translates to a seed length of $n/2 + 2\log(1/\varepsilon) + O(1)$. There is a lot of work on the problems of proving existence of and explicitly constructing Ramanujan graphs and “near-Ramanujan” graphs [\[LPS88; Mar88; Mor94; Fri08; CM08; BT11; MSS15; Coh16; Bor20; MOP20; Alo21\]](#).

3.1.2 Combinatorial rectangles and the Expander Mixing Lemma

In the language of PRGs, the famous Expander Mixing Lemma is the “1 \implies 2” direction of [Lemma 3.1.4](#), specialized to the case that g and h are Boolean-valued. That is, the Expander Mixing Lemma explains how to use expander graphs to fool *two-dimensional combinatorial rectangles*, defined next.

Definition 3.1.6 (Combinatorial rectangles). *Let n be a multiple of d and let $f: \{0, 1\}^n \rightarrow \{0, 1\}$. We say that f is a d -dimensional combinatorial rectangle if there are functions $f_1, \dots, f_d: \{0, 1\}^{n/d} \rightarrow \{0, 1\}$ such that*

$$f(x^{(1)}, \dots, x^{(d)}) = \prod_{i=1}^d f_i(x^{(i)}).$$

Lemma 3.1.7 (Expander Mixing Lemma). *Let n be an even positive integer and let G be an ε -spectral expander on vertex set $\{0, 1\}^{n/2}$. Sample a uniform random vertex X , then sample a uniform random neighbor Y of X . Then the distribution (X, Y) fools two-dimensional combinatorial rectangles with error ε .*

Proof. Immediate from [Lemma 3.1.4](#). □

Corollary 3.1.8 (Optimal PRG for two-dimensional combinatorial rectangles). *For every $n \in \mathbb{N}$ and every $\varepsilon > 0$, there is an explicit ε -PRG for two-dimensional combinatorial rectangles on n bits with seed length $\frac{n}{2} + O(\log(1/\varepsilon))$.*

Proof. Immediate from [Lemma 3.1.7](#) and [Theorem 3.1.5](#). □

There is a large body of work designing PRGs for high-dimensional combinatorial rectangles [[ASWZ96](#); [EGLNV98](#); [LLSZ97](#); [Lu02](#); [GMRTV12](#); [Vio14](#); [GY20](#); [HLV18](#); [Lee19](#)]. Near-optimal constructions are known [[GMRTV12](#); [GY20](#); [Lee19](#)], but it is still an open problem to get the optimal seed length.

Open Problem 3.1.9 (Optimal PRGs for high-dimensional combinatorial rectangles). *Design an explicit PRG for d -dimensional combinatorial rectangles with seed length $O(n/d + \log(1/\varepsilon) + \log \log n)$.*

For this section, however, it suffices to focus on the two-dimensional case. A two-dimensional combinatorial rectangle f can be computed by a communication protocol with a particularly simple structure: Alice computes $g(x)$ and sends it to Bob, and then Bob computes $h(y)$ and multiplies. Given [Corollary 3.1.8](#), we are now ready to fool *general* two-party communication protocols by using the standard decomposition of communication protocols into combinatorial rectangles.

Proof of Theorem 3.1.2. Without loss of generality, we may assume that the parties always transmit exactly m bits in total. Define $\pi: \{0, 1\}^n \rightarrow \{0, 1\}^m$ by letting $\pi(x, y)$ be the sequence of bits transmitted when Alice holds x and Bob holds y . For each possible transcript $z \in \{0, 1\}^m$, define functions $g_z, h_z: \{0, 1\}^{n/2} \rightarrow \{0, 1\}$ by the rule

$$\begin{aligned} g_z(x) = 1 &\iff \exists y, \pi(x, y) = z \\ h_z(y) = 1 &\iff \exists x, \pi(x, y) = z. \end{aligned}$$

That way, $\pi(x, y) = z \iff g_z(x) = h_z(y) = 1$. Since both parties output $f(x, y)$, there is some set $A \subseteq \{0, 1\}^m$ such that $f(x, y) = 1 \iff \pi(x, y) \in A$. Therefore,

$$f(x, y) = \sum_{z \in A} g_z(x) \cdot h_z(y),$$




Figure 3.1: A width-5 length-6 standard-order ROBP computing the function $f(x) = \text{MAJ}(x_1 \oplus x_2, x_3 \oplus x_4, x_5 \oplus x_6)$.

a linear combination of two-dimensional combinatorial rectangles. Therefore, by the Triangle Inequality for PRG Errors, every δ -PRG for two-dimensional combinatorial rectangles fools f with error $|A|\delta \leq 2^m\delta$. Picking $\delta = 2^{-m}\varepsilon$ and applying [Corollary 3.1.8](#) completes the proof. \square

3.2 The INW generator for standard-order ROBPs

In this section, we will present the Impagliazzo-Nisan-Wigderson (INW) generator [[INW94](#)], which is one of the most influential unconditional PRG constructions. We will prove that it fools *standard-order read-once branching programs* (standard-order ROBPs), defined next.

Definition 3.2.1 (Standard-order read-once branching programs). *A length- n standard-order read-once branching program (standard-order ROBP) f consists of a directed layered multigraph with $n + 1$ layers, V_0, \dots, V_n . For every $i < n$, each vertex $v \in V_i$ has two outgoing edges leading to V_{i+1} , one labeled 0 and the other labeled 1. Vertices in V_n have zero outgoing edges. There is a designated “start vertex” $v_{\text{start}} \in V_0$. An input $x \in \{0, 1\}^n$ selects a path (v_0, v_1, \dots, v_n) through the graph: the path starts at $v_0 = v_{\text{start}}$, and upon reaching a vertex $v_i \in V_i$, the bit x_{i+1} specifies which outgoing edge to use. There is a designated set of “accept vertices” $V_{\text{accept}} \subseteq V_n$, and $f(x) = 1$ if $v_n \in V_{\text{accept}}$ and $f(x) = 0$ otherwise. The width of the program is the maximum number of vertices in a single layer. (See [Figure 3.1](#).)*

The standard-order ROBP model is important because as discussed in [Section 1.1.3](#), it models the behavior of a space-efficient randomized algorithm as a function of its random bits. In more detail, if σ is an input and A is a randomized space- S decision algorithm where $S \geq \log |\sigma|$, then there is some width- w ROBP f computing the function $f(x) = A(\sigma, x)$, where $w = 2^{O(S)}$. Therefore,

if X is a distribution that fools width- w ROBPs, then A can be executed using X in place of true randomness without significantly distorting the algorithm’s acceptance probability.

Remark 3.2.2 (ROBP terminology). *In the pseudorandomness literature, standard-order ROBPs are often referred to as simply “ROBPs.” This practice is a bit misleading, since the definition is not simply “a branching program that is read-once.” Indeed, in addition to being read-once, we are assuming that the program is oblivious, meaning that the variable queried in time step i depends only on i , and more specifically, we are assuming that the branching program follows the standard variable ordering, meaning that in time step i , the program queries the variable x_i . (These assumptions are well-motivated, because in the standard model of randomized space-bounded computation, we merely have the ability to repeatedly flip a single fair coin, which is equivalent to observing a one-way stream of random bits.) Unsurprisingly, many papers outside the pseudorandomness literature use terms like “read-once branching program” to refer to more general models that are not necessarily even oblivious [Weg87; BHST87; Raz91; BJS01; BLP15]. In this text, for clarity, we use the more verbose term “standard-order ROBP” to emphasize the variable ordering assumption.¹*

There is some convenient standard notation for *subprograms* of ROBPs.

Definition 3.2.3 (Subprograms of ROBPs). *Let f be a standard-order ROBP with layers V_0, \dots, V_n . Let $u \in V_i$ and $S \subseteq V_j$ where $i \leq j$. We let $f_{u \rightarrow S}$ denote the subprogram from u to S , i.e., the standard-order ROBP on layers V_i, \dots, V_j with start vertex u and accept vertices S . We use the shorthand $f_{u \rightarrow v} = f_{u \rightarrow \{v\}}$, $f_{\rightarrow v} = f_{v_{\text{start}} \rightarrow v}$, and $f_{u \rightarrow} = f_{u \rightarrow V_{\text{accept}}}$.*

3.2.1 Concatenating two independent pseudorandom strings

As mentioned previously, we are working toward presenting the INW PRG and proving that it fools standard-order ROBPs. The INW PRG is based on recycling seeds *recursively*. After constructing a PRG with output length $n/2$, we will use two correlated seeds to generate two pseudorandom strings of length $n/2$ and concatenate them to get a string of length n . To argue that the INW PRG works, we will argue that using two correlated seeds is almost as good as using two independent seeds.

Let us begin by arguing that using two independent seeds would work well in the first place. This argument is generic and applies to many different models of computation, not just ROBPs.

Lemma 3.2.4 (Concatenating independent pseudorandom strings). *Let n be an even positive integer, let $f: \{0, 1\}^n \rightarrow \mathbb{R}$, and let X_L, X_R be independent random variables distributed over $\{0, 1\}^{n/2}$. Let \mathcal{F} be the class of all restrictions of f to $n/2$ bits. If X_L and X_R fool \mathcal{F} with error ε , then (X_L, X_R) fools f with error 2ε .*

Proof. Sample $U \sim U_{n/2}$ so that X_L, X_R, U are mutually independent. Then

$$\begin{aligned} \mathbb{E}_{X_L, X_R} [f(X_L, X_R)] &= \mathbb{E}_{X_L} \left[\mathbb{E}_{X_R} [f(X_L, X_R)] \right] = \mathbb{E}_{X_L} \left[\mathbb{E}_U [f(X_L, U)] \pm \varepsilon \right] \\ &= \mathbb{E}_U \left[\mathbb{E}_{X_L} [f(X_L, U)] \right] \pm \varepsilon \\ &= \mathbb{E}[f] \pm 2\varepsilon. \quad \square \end{aligned}$$

¹Hoza used the same verbose terminology in some other recent expository work [Hoz22].

3.2.2 Recycling seeds using a PRG for two-dimensional rectangles

The main lemma of the INW generator allows us to double the output length of a PRG by paying a small *additive* penalty in terms of the seed length.

Lemma 3.2.5 (Recycling seeds for standard-order ROBPs). *Let n be an even positive integer. Suppose $G: \{0, 1\}^s \rightarrow \{0, 1\}^{n/2}$ is an ε -PRG for width- w length- $(n/2)$ standard-order ROBPs. Let (Y_L, Y_R) be a distribution over $\{0, 1\}^{2s}$ that δ -fools two-dimensional combinatorial rectangles. Then $(G(Y_L), G(Y_R))$ fools width- w length- n standard-order ROBPs with error $2\varepsilon + w\delta$.*

Proof. Let f be a width- w length- n standard-order ROBP. Define $g: \{0, 1\}^{2s} \rightarrow \{0, 1\}$ by composing f with two independent copies of G , i.e.,

$$g(y_L, y_R) = f(G(y_L), G(y_R)).$$

By Lemma 3.2.4, using two independent seeds fools f with error 2ε , i.e., $|\mathbb{E}[f] - \mathbb{E}[g]| \leq 2\varepsilon$. Now let us compare two independent seeds to the two correlated seeds (Y_L, Y_R) . For any $x_L, x_R \in \{0, 1\}^{n/2}$, we can write

$$f(x_L, x_R) = \sum_{v \in V_{n/2}} f_{\rightarrow v}(x_L) \cdot f_{v \rightarrow}(x_R).$$

Consequently, if we define $g_v(y_L, y_R) = f_{\rightarrow v}(G(y_L)) \cdot f_{v \rightarrow}(G(y_R))$, then

$$g(y_L, y_R) = \sum_{v \in V_{n/2}} g_v(y_L, y_R).$$

Each function g_v is a two-dimensional combinatorial rectangle on $2s$ bits. Therefore, (Y_L, Y_R) fools g_v with error δ , so by the Triangle Inequality for PRG Errors, (Y_L, Y_R) fools g with error $w\delta$. Therefore,

$$|\mathbb{E}[f] - \mathbb{E}[g(Y_L, Y_R)]| \leq 2\varepsilon + w\delta. \quad \square$$

Theorem 3.2.6 (The INW generator for standard-order ROBPs [INW94]). *For every $n, w \in \mathbb{N}$ and every $\varepsilon > 0$, there is an explicit ε -PRG for width- w length- n standard-order ROBPs with seed length $O(\log(wn/\varepsilon) \cdot \log n)$.*

Proof. Set $\delta = \frac{\varepsilon}{wn}$. We start with the trivial PRG $G(y) = y$ where $|y| = 1$. We repeatedly apply Lemma 3.2.5, using Corollary 3.1.8 to sample Y_L, Y_R . Each time we apply Lemma 3.2.5, we double the output length of our PRG, we increase its seed length from s to $s + O(\log(1/\delta)) = s + O(\log(wn/\varepsilon))$, and we increase its error from ε_0 to $2\varepsilon_0 + \frac{\varepsilon}{n}$. After $\log n$ iterations, we are done. \square

The seed length in Theorem 3.2.6 was actually already achieved by Nisan [Nis92] prior to Impagliazzo, Nisan, and Wigderson's work [INW94]. Nisan's PRG [Nis92] follows a fairly similar intuition as the INW generator, but the details are different. The INW generator has some advantages over Nisan's generator; most importantly, the INW generator has turned out to be more flexible and amenable to analysis in other models. (We will see an example in Section 3.3.)

The optimal seed length for fooling width- w length- n standard-order ROBPs would be $O(\log(wn/\varepsilon))$. Designing optimal or near-optimal PRGs for standard-order ROBPs is one of the biggest open problems in the unconditional theory of PRGs. Some of the work on this problem focuses on the case that the width of the program is very small. In Section 2.3.4, we saw that small-bias generators fool width-2 branching programs with seed length $O(\log(n/\varepsilon))$. Explicit PRGs for width-3 standard-order ROBPs are known with seed length $\tilde{O}(\log n \cdot \log(1/\varepsilon))$ [MRT19]. However, for width-4 ROBPs, no PRG constructions are known with a seed length better than that of the INW generator.

Open Problem 3.2.7 (Better PRGs for width-4 ROBPs). *Design an explicit 0.1-PRG for width-4 length- n standard-order ROBPs with seed length $o(\log^2 n)$.*

A candidate PRG for ROBPs, suggested by Reingold and Vadhan [MZ09; LV17], is to take a sum (i.e., bitwise XOR) of a few independent small-bias distributions. Recall Viola’s proof that a sum of d small-bias distributions fools degree- d polynomials over \mathbb{F}_2 (see Section 2.4). Perhaps a similar PRG can work for constant-width ROBPs.

Open Problem 3.2.8 (ROBPs and sums of small-bias distributions). *Prove or disprove the suggestion that a sum of $O(1)$ small-bias distributions fools constant-width standard-order ROBPs with near-logarithmic seed length.*

Amazingly, it is consistent with current knowledge that simply summing two small-bias distributions fools *polynomial*-width standard-order ROBPs with optimal seed length $O(\log(n/\varepsilon))$.

Due to the difficulty of constructing improved PRGs, much of the recent research on pseudorandomness for ROBPs has focused on constructing relaxations of PRGs such as HSGs and WPRGs (defined in Section 1.5) [GMRTV12; ŠZ21; BCG20; HZ20; CL20; CDRST21; PV21b; Hoz21]. Another line of work seeks PRGs for *restricted classes* of ROBPs; we will discuss an example in the next section.

3.3 The BRRY generator for standard-order regular ROBPs

As discussed in the previous section, it is still an open problem to design an explicit PRG for constant-width standard-order ROBPs with seed length $o(\log^2 n)$. However, we can get a better seed length for many interesting special cases, including *regular* programs. A standard-order ROBP is called *regular* if every vertex of the program (except those in the first layer) has in-degree 2; see Figure 3.2. There is a reduction showing that optimal PRGs for standard-order regular ROBPs imply optimal PRGs for *all* standard-order ROBPs [RTV06; BHPP22], which helps to motivate studying this special case. We will present a result by Braverman, Rao, Raz, and Yehudayoff [BRRY14], which shows how to achieve seed length $\tilde{O}(\log n)$ for constant-width standard-order regular ROBPs.

Theorem 3.3.1 (The BRRY generator for standard-order regular ROBPs [BRRY14]). *For every $w, n \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for width- w length- n standard-order regular ROBPs with seed length*

$$O(\log n \cdot (\log w + \log \log n + \log(1/\varepsilon))).$$

As we will see, the PRG *construction* is actually the same INW generator that we discussed in Section 3.2 (albeit with different parameters). The improvement comes from a better analysis.

3.3.1 Improved analysis of the INW generator for low-weight programs

For intuition, let us briefly summarize the INW analysis that we seek to improve (see Section 3.2). In each round of the INW construction, we use a δ -spectral expander to recycle the seed from the previous round. The analysis essentially argues that we pay an error δ at each vertex in the program, so the total error is δwn . We therefore chose $\delta = \frac{\varepsilon}{wn}$. In each of the $\log n$ rounds of the construction, the seed length increases by an additive $O(\log(1/\delta))$, leading to the final seed length of $O(\log(wn/\varepsilon) \cdot \log n)$.

Braverman, Rao, Raz, and Yehudayoff had the insight that this analysis is overly pessimistic in some cases, because it treats every vertex in the program as “important.” In reality, sometimes a vertex v is “unimportant,” in the sense that its two out-neighbors $v[0], v[1]$ have almost the same




Figure 3.2: Let n be an even positive integer. The function $f(x) = x_1x_2 \oplus x_3x_4 \oplus \cdots \oplus x_{n-1}x_n$ can be computed by a width-4 standard-order regular ROBP. The case $n = 4$ is shown above.

acceptance probabilities: $\mathbb{E}[f_{v[0]\rightarrow}] \approx \mathbb{E}[f_{v[1]\rightarrow}]$. After reaching such a v , it doesn't matter much whether we read a high-quality random bit or a low-quality random bit, because it doesn't matter much whether we go to $v[0]$ or $v[1]$. Rather than contributing a penalty of δ to the overall error bound, intuitively we might hope that the error at such a vertex is closer to $\delta \cdot |\mathbb{E}[f_{v[0]\rightarrow}] - \mathbb{E}[f_{v[1]\rightarrow}]|$.

To formalize this intuition, we rely on a generalization of ROBPs that output real values instead of Boolean ones, called *read-once evaluation programs* (ROEPs).

Definition 3.3.2 (Read-once evaluation programs). *A length- n standard-order read once evaluation program (standard-order ROEP) f has the same graph structure as a length- n standard-order ROBP, but in the final layer, instead of a set of accept vertices $V_{\text{accept}} \subseteq V_n$, it has a value $q_v \in \mathbb{R}$ assigned to each vertex $v \in V_n$. Each input $x \in \{0, 1\}^n$ defines a path from v_{start} to a vertex $v \in V_n$ as usual, which determines the output of the program $f(x) = q_v$. Thus, f computes a function $f: \{0, 1\}^n \rightarrow \mathbb{R}$.*

Standard-order ROBPs are a special case of standard-order ROEPs, where $q_v = 1$ if $v \in V_{\text{accept}}$ and $q_v = 0$ if $v \in V_n \setminus V_{\text{accept}}$. If v is a vertex in an ROEP f , then the subprogram $f_{\rightarrow v}$ is an ROBP while the subprogram $f_{v \rightarrow}$ is an ROEP. We extend the notation q_v to the case that $v \notin V_n$ by the rule $q_v = \mathbb{E}[f_{v \rightarrow}]$, i.e., q_v is the expected label of the vertex reached by starting at v and taking a random walk to V_n . The ‘‘importance’’ of a vertex in an ROEP is captured by the following definition.

Definition 3.3.3 (Weight in an ROEP). *Let f be a standard-order ROEP with layers V_0, \dots, V_n . Let $v \in V_i$ with $i < n$. Let $v[0]$ and $v[1]$ be the two out-neighbors of v . The weight of v , denoted $\text{Weight}(v)$, is defined to be*

$$\text{Weight}(v) = |q_{v[0]} - q_{v[1]}|.$$

The weight of f is the sum of the weights of the vertices,² i.e.,

$$\text{Weight}(f) = \sum_{i=0}^{n-1} \sum_{v \in V_i} \text{Weight}(v).$$

²In the original paper [BRRY14], Braverman, Rao, Raz, and Yehudayoff considered edge weights instead of vertex weights. The two definitions are equivalent.

In Fourier-analytic terms, we have the following formula [RSV13]:

$$\text{Weight}(f) = 2 \sum_{i=1}^n \sum_{v \in V_{i-1}} \left| \widehat{f_{v \rightarrow \{i\}}} \right|.$$

Clearly, we always have $\text{Weight}(v) \leq 1$ and hence $\text{Weight}(f) \leq wn$. Now let us show that the INW generator fools low-weight programs with a shorter seed.

Theorem 3.3.4 (The BRRY generator for low-weight standard-order ROEPs [BRRY14]). *For every $w, n, m \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for width- w length- n ROEPs f that satisfy $\text{Weight}(f) \leq m$ with seed length*

$$O(\log n \cdot (\log(wm/\varepsilon) + \log \log n)).$$

In the original paper [BRRY14], to prove [Theorem 3.3.4](#), Braverman, Rao, Raz, and Yehudayoff analyzed the INW generator in terms of *randomness extraction*, similar to the analysis of the Nisan-Zuckerman generator (see [Section 3.4](#)). Here, we will show how to carry out the analysis more directly, using the PRG characterization of spectral expanders ([Lemma 3.1.4](#)). The first step, like the original INW analysis, is to analyze two independent seeds. We use the following refinement of [Lemma 3.2.4](#).

Lemma 3.3.5 (Refined analysis of the concatenation of independent pseudorandom strings). *Let n be an even positive integer, let $f: \{0, 1\}^n \rightarrow \mathbb{R}$, and let X_L, X_R be independent random variables distributed over $\{0, 1\}^{n/2}$. Define $f_L: \{0, 1\}^{n/2} \rightarrow \mathbb{R}$ by*

$$f_L(x) = \mathbb{E}[f(x, U_{n/2})],$$

and let \mathcal{F}_R be the class of all functions $f_R: \{0, 1\}^{n/2} \rightarrow \mathbb{R}$ of the form

$$f_R(x) = f(a, x)$$

where $a \in \{0, 1\}^{n/2}$. If X_L fools f_L with error ε_L and X_R fools \mathcal{F}_R with error ε_R , then (X_L, X_R) fools f with error $\varepsilon_L + \varepsilon_R$.

Proof. Sample $U \sim U_{n/2}$ so that X_L, X_R, U are mutually independent. Then

$$\begin{aligned} \mathbb{E}_{X_L, X_R} [f(X_L, X_R)] &= \mathbb{E}_{X_L} \left[\mathbb{E}_{X_R} [f(X_L, X_R)] \right] = \mathbb{E}_{X_L} \left[\mathbb{E}_U [f(X_L, U)] \pm \varepsilon_R \right] \\ &= \mathbb{E}_{X_L} [f_L(X_L)] \pm \varepsilon_R \\ &= \mathbb{E}[f_L] \pm (\varepsilon_L + \varepsilon_R) \\ &= \mathbb{E}[f] \pm (\varepsilon_L + \varepsilon_R). \quad \square \end{aligned}$$

Next, we argue that the outputs of a very-low-weight ROEP fall in a small interval.

Lemma 3.3.6 (Low weight \implies bounded range). *For every standard-order ROEP f and every pair of inputs x, y , we have $|f(x) - f(y)| \leq \text{Weight}(f)$.*

Proof. We proceed by induction on the length of f . For a length-1 standard-order ROEP, by definition, $|f(0) - f(1)| = \text{Weight}(v_{\text{start}}) \leq \text{Weight}(f)$. If the length of f is $n > 1$, then let f_L be the ROEP consisting of the first $n - 1$ layers of f , and let $u, v \in V_{n-1}$ be the penultimate vertices visited when f reads x and y respectively. By induction, $|q_u - q_v| \leq \text{Weight}(f_L)$. Furthermore, we have $|q_u - f(x)| = \frac{1}{2} \text{Weight}(u)$ and $|q_v - f(y)| = \frac{1}{2} \text{Weight}(v)$. Therefore,

$$|f(x) - f(y)| \leq \text{Weight}(f_L) + \frac{1}{2} \text{Weight}(u) + \frac{1}{2} \text{Weight}(v) \leq \text{Weight}(f). \quad \square$$

Finally, we are ready to analyze two correlated seeds. Recall that the INW generator is based on a key lemma ([Lemma 3.2.5](#)) that says, if G fools width- w length- $(n/2)$ ROBPs with error ε and (X, Y) is a random edge in a δ -spectral expander, then $(G(X), G(Y))$ fools width- w length- n ROBPs with error $2\varepsilon + \delta w$. We will prove the following more refined lemma that allows us to avoid the critical factor of two.

Lemma 3.3.7 (Recycling randomness for low-weight programs). *Suppose $G: \{0, 1\}^s \rightarrow \{0, 1\}^{n/2}$ fools every width- w length- $(n/2)$ ROEP f with error $\varepsilon \cdot \text{Weight}(f)$. Fix some δ -spectral expander on the vertex set $\{0, 1\}^s$, and sample a uniform random vertex X and a uniform random neighbor Y . Then the distribution $(G(X), G(Y))$ fools every width- w length- n ROEP f with error $(\varepsilon + w\delta) \cdot \text{Weight}(f)$.*

Proof. Let f be a width- w length- n ROEP. Let Weight_L and Weight_R denote the weights of the left half and right half of f , respectively, so $\text{Weight}(f) = \text{Weight}_L + \text{Weight}_R$. Similarly to the proof of [Lemma 3.2.5](#), we can write

$$f(x, y) = \sum_{v \in V_{n/2}} f_{\rightarrow v}(x) \cdot f_{v \rightarrow}(y).$$

Fix a vertex $v \in V_{n/2}$. Define $g(a) = f_{\rightarrow v}(G(a)) \in \{0, 1\}$ and $h(b) = f_{v \rightarrow}(G(b)) \in \mathbb{R}$. By [Lemma 3.3.6](#), the outputs of h fall in an interval of length $\text{Weight}(f_{v \rightarrow}) = \text{Weight}_R$. Therefore, by [Lemma 3.1.4](#), we have

$$|\mathbb{E}[g(X) \cdot h(Y)] - \mathbb{E}[g] \mathbb{E}[h]| \leq \delta \cdot \text{Var}[g] \cdot \text{Var}[h] \leq \delta \cdot 1 \cdot \text{Weight}_R.$$

Therefore, if we let U, U' be independent uniform seeds, then by the triangle inequality,

$$|\mathbb{E}[f(G(U), G(U')))] - \mathbb{E}[f(G(X), G(Y))]| \leq \sum_{v \in V_{n/2}} \delta \cdot \text{Weight}_R = w\delta \cdot \text{Weight}_R.$$

To bound $|\mathbb{E}[f(G(U), G(U')))] - \mathbb{E}[f]|$, we use [Lemma 3.3.5](#). The function f_L that appears in that lemma is precisely the left half of f , a standard-order ROEP of weight Weight_L . Therefore, G fools f_L with error $\varepsilon \cdot \text{Weight}_L$. Meanwhile, each function f_R considered in that lemma is a subprogram of the form $f_{v \rightarrow}$ for some $v \in V_{n/2}$, hence a standard-order ROEP of weight Weight_R . Therefore, G fools f_R with error $\varepsilon \cdot \text{Weight}_R$. Thus, [Lemma 3.3.5](#) guarantees that

$$|\mathbb{E}[f(G(U), G(U')))] - \mathbb{E}[f]| \leq \varepsilon \cdot \text{Weight}_L + \varepsilon \cdot \text{Weight}_R.$$

Therefore,

$$|\mathbb{E}[f(G(X), G(Y)))] - \mathbb{E}[f]| \leq \varepsilon \cdot \text{Weight}(f) + w\delta \cdot \text{Weight}_R \leq (\varepsilon + w\delta) \cdot \text{Weight}(f). \quad \square$$

Just like the original INW generator, we can use [Lemma 3.3.7](#) to inductively construct a PRG for width- w length- n standard-order ROEPs. We start with a trivial PRG outputting a single bit. Then, in each of $\log n$ steps, we use a δ -spectral expander to recycle the seed and thereby double the output length. The final error is $w\delta \cdot \text{Weight}(f) \cdot \log n$. To ε -fool standard-order ROEPs of weight at most m , we may set $\delta = \frac{\varepsilon}{wm \log n}$. If we use a sparse expander (see [Theorem 3.1.5](#)), then in each step, the seed length of our PRG increases by an additive $O(\log(1/\delta))$ bits. Thus, the final seed length is $O(\log n \cdot \log(1/\delta))$, which is

$$O(\log n \cdot (\log(wm/\varepsilon) + \log \log n)),$$

completing the proof of [Theorem 3.3.4](#).

3.3.2 Regular programs have low weight

Recall that our original goal was to design an improved PRG for standard-order regular ROBPs (Theorem 3.3.1). To apply the analysis of low-weight programs, Braverman, Rao, Raz, and Yehudayoff showed that width- w standard-order regular ROBPs have weight at most $O(w^2)$.³ The original proof has a “combinatorial” flavor and is based on studying a certain pebble game. Here we present a more “analytic” interpretation of their argument.

Lemma 3.3.8 (Regular programs have bounded weight). *If f is a length- n standard-order regular ROEP, then*

$$\text{Weight}(f) \leq 2 \sum_{u,v \in V_n} |q_u - q_v|.$$

In particular, if f is a width- w standard-order regular ROBP, then $\text{Weight}(f) \leq 2w^2$.

Proof. For each i , let $D_i = \sum_{u,v \in V_i} |q_u - q_v|$. For $i < n$, let $\text{Weight}_i = \sum_{u \in V_i} \text{Weight}(u)$. For a vertex u , we let $u[0]$ and $u[1]$ denote the two out-neighbors of u . If we sample $X, Y \in \{0, 1\}$ uniformly at random, then

$$\begin{aligned} \frac{1}{2} \text{Weight}_i + D_i &= \frac{1}{2} \sum_{u \in V_i} |q_{u[0]} - q_{u[1]}| + \sum_{\substack{u,v \in V_i \\ u \neq v}} |q_u - q_v| \\ &= \sum_{u \in V_i} \mathbb{E}_{X,Y} [|q_{u[X]} - q_{u[Y]}|] + \sum_{\substack{u,v \in V_i \\ u \neq v}} \left| \mathbb{E}_X [q_{u[X]}] - \mathbb{E}_Y [q_{v[Y]}] \right| \\ &\leq \sum_{u \in V_i} \mathbb{E}_{X,Y} [|q_{u[X]} - q_{u[Y]}|] + \sum_{\substack{u,v \in V_i \\ u \neq v}} \mathbb{E}_{X,Y} [|q_{u[X]} - q_{v[Y]}|] \quad (\text{Triangle inequality}) \\ &= \sum_{u,v \in V_i} \mathbb{E}_{X,Y} [|q_{u[X]} - q_{v[Y]}|] \\ &= \sum_{u,v \in V_{i+1}} \frac{\deg^-(u) \cdot \deg^-(v)}{4} \cdot |q_u - q_v| \\ &= D_{i+1} \quad (\text{Regularity.}) \end{aligned}$$

Rearranging, we have shown $\frac{1}{2} \text{Weight}_i \leq D_{i+1} - D_i$. Summing over $i < n$ completes the proof. \square

Combining Lemma 3.3.8 and Theorem 3.3.4 completes the proof of Theorem 3.3.1.

Besides regular ROBPs, another well-studied class of branching programs is *permutation ROBPs*, which are regular ROBPs with the additional assumption that for each vertex v , the two incoming edges of v have distinct labels. Braverman, Rao, Raz, and Yehudayoff’s paper [BRRY14] is one of a great number of papers studying pseudorandomness for regular and permutation ROBPs [BRRY14; BV10; De11; KNP11; Ste12; RSV13; CHHL19; HPV21; PV21b; PV21a; PV22; BHPP22; LPV22]. We will revisit permutation ROBPs in Section 5.2.

Let us highlight one open problem regarding these topics. As mentioned in Section 3.2, Meka, Reingold, and Tal designed a PRG for width-3 standard-order ROBPs with seed length $\tilde{O}(\log n \cdot$

³In Braverman, Rao, Raz, and Yehudayoff’s definition of a regular ROBP [BRRY14], they only allowed a single accept vertex ($|V_{\text{accept}}| = 1$), hence their weight bound was $O(w)$ rather than $O(w^2)$. For the purpose of Theorem 3.3.1, it makes no difference whether we allow multiple accept vertices, because an ε -PRG for the single-accept-vertex model is also an (εw) -PRG for the multiple-accept-vertex model. However, there are other contexts in which bounding the number of accept vertices makes a huge difference [HPV21; PV21b; PV21a; PV22; BHPP22; LPV22].

$\log(1/\varepsilon)$) [MRT19]. Roughly speaking, their approach involves reducing to the permutation case and then applying the INW generator. The INW generator has a $\log n \cdot \log(1/\varepsilon)$ term in its seed length (even for permutation branching programs [PV21a]), so Meka, Reingold, and Tal get no improvement over Nisan’s PRG when the error is $1/n$. This motivates the following problem.

Open Problem 3.3.9 (Low-error PRGs for width-3 standard-order permutation ROBPs). *Design an explicit PRG for width-3 standard-order permutation ROBPs with error $1/n$ and seed length $o(\log^2 n)$.*

3.4 The Nisan-Zuckerman generator for short, wide ROBPs

In the previous section, we focused on width- w length- n standard-order ROBPs where $w \ll n$. In this section, let us study the opposite regime, i.e., $w \gg n$. One reason to study this regime is that it corresponds to derandomizing space-bounded algorithms that only use a little bit of randomness in the first place.

Nisan and Zuckerman designed a PRG with *optimal* seed length $O(\log w)$ for the case that $n = \text{polylog } w$ and the error parameter is moderate [NZ96]. In contrast, the INW generator’s seed length is $\Theta(\log w \cdot \log \log w)$ for such parameters.

Theorem 3.4.1 (The Nisan-Zuckerman generator). *Let $\nu > 0$ and $c \geq 1$ be constants. For all $w \in \mathbb{N}$, there is an explicit PRG for width- w length- $(\log^c w)$ standard-order ROBPs with seed length $O(\log w)$ and error $2^{-\log^{1-\nu} w}$.*

Recall that when we say the PRG is “explicit,” we mean that its time complexity is polynomial in its output length (Definition 1.4.1). In the case of Theorem 3.4.1, this means that the time complexity of the PRG is $\text{polylog}(w)$. As usual in this text, we will refrain from carefully verifying this time complexity bound. However, it will turn out to be useful to analyze the *space complexity* of the PRG. As we will see, the Nisan-Zuckerman PRG can be computed using $O(\log w)$ bits of space. Using the connection between randomized space-bounded algorithms and ROBPs, one can show the following striking corollary: If a decision problem can be solved by a randomized algorithm using S bits of space and $\text{poly}(S)$ random bits, then it can also be solved by a deterministic algorithm using $O(S)$ bits of space [NZ96].

3.4.1 Randomness extractors

The proof of Theorem 3.4.1 uses seeded *randomness extractors* (introduced by Nisan and Zuckerman [NZ96]) to recycle randomness. Informally, a randomness extractor is a tool for converting an “imperfect” source of randomness into near-uniform random bits. The following definition specifies our model of imperfect randomness.

Definition 3.4.2 (k -Source [CG88; Zuc90]). *The min-entropy of a distribution X , denoted $H_{\min}(X)$, is defined by*

$$H_{\min}(X) = \inf_{x \in \text{supp}(X)} \log \left(\frac{1}{\Pr[X = x]} \right).$$

We say that a random variable X on $\{0, 1\}^t$ is a k -source if $H_{\min}(X) \geq k$, that is, for every $x \in \{0, 1\}^t$, $\Pr[X = x] \leq 2^{-k}$.

(Min-entropy is also sometimes denoted $H_{\infty}(X)$.) To precisely define extractors, let $d_{\text{TV}}(\cdot, \cdot)$ denote total variation distance.

Definition 3.4.3 (Seeded Randomness Extractor). *A function $\text{Ext}: \{0, 1\}^t \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε) -extractor if for every k -source X ,*

$$d_{\text{TV}}(\text{Ext}(X, U_d), U_m) \leq \varepsilon,$$

where U_d is independent of X .

Intuitively, extractors and PRGs both produce some type of random bits, but they are incomparable:

- The output of an extractor fools *all* tests $f: \{0, 1\}^m \rightarrow \{0, 1\}$ (this is equivalent to being close to uniform in total variation distance), whereas the output of a PRG only fools *some* tests.
- On the other hand, an extractor requires a seed *and* an imperfect source of randomness, whereas a PRG only requires a seed.

Formally, it is possible to view extractors as a special case of PRGs,⁴ but this viewpoint is probably not helpful.

There is a rich and beautiful theory of randomness extractors that goes beyond the scope of this text; see, e.g., Vadhan’s monograph [Vad12]. We will take for granted an explicit construction due to Guruswami, Umans, and Vadhan [GUV09], or rather a space-optimized version by Kane, Nelson, and Woodruff [KNW08].

Theorem 3.4.4 ([GUV09; KNW08]). *For every $k \leq t$ and $\varepsilon > 0$, there is an (k, ε) -extractor $\text{Ext}: \{0, 1\}^t \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $m \geq k/2$ and $d = O(\log(t/\varepsilon))$. Furthermore, given (x, y, k) as input, $\text{Ext}(x, y)$ can be computed in space $O(t + \log(1/\varepsilon))$.*

Toward proving [Theorem 3.4.1](#), we will first study a PRG of the following form:

$$G(x, y_1, \dots, y_\ell) = (\text{Ext}(x, y_1), \dots, \text{Ext}(x, y_\ell)). \quad (3.1)$$

We will then compose this PRG with itself to prove [Theorem 3.4.1](#). Here’s the intuition behind [Equation \(3.1\)](#). Let (X, Y_1, \dots, Y_ℓ) be a uniform random seed. After the ROBP reads some prefix $(\text{Ext}(X, Y_1), \dots, \text{Ext}(X, Y_i))$, it only “remembers” $\log w$ bits of information about what it has seen. We will set $|X| = 3 \log w$. Since the ROBP only “knows” $\log w$ bits about X , the random variable X should still have $2 \log w$ bits of entropy “from the ROBP’s perspective.” Therefore, $\text{Ext}(X, Y_{i+1})$ should appear nearly uniform to the ROBP.

The most elegant way to formalize this intuition is to use the concept of conditional min-entropy introduced by Dodis, Ostrovsky, Reyzin, and Smith [DORS08].

Definition 3.4.5 (Conditional min-entropy). *Let X and A be jointly distributed random variables. The conditional min-entropy of X given A is*

$$\tilde{H}_{\min}(X | A) = \log \left(\frac{1}{\mathbb{E}_{a \sim A} [\sup_{x \in \text{supp}(X)} \Pr[X = x | A = a]]} \right).$$

Conditional min-entropy can be interpreted in terms of strategies for guessing X after seeing A [Vad12, Problem 6.7]. Conditional min-entropy satisfies the following intuitive “chain rule” first proven by Dodis, Ostrovsky, Reyzin, and Smith [DORS08, Lemma 2.2].

⁴Given $\text{Ext}: \{0, 1\}^t \times \{0, 1\}^d \rightarrow \{0, 1\}^m$, we can define $G: \{0, 1\}^d \rightarrow \{0, 1\}^{2^t \times m}$ by letting $G(y)$ be the table of all values $\text{Ext}(x, y)$ for $x \in \{0, 1\}^t$. Then Ext is a (k, ε) -extractor if and only if G is an ε -PRG for the class of functions f of the form $f(M) = \frac{1}{|S|} \cdot \sum_{i \in S} g(M_i)$, where $g: \{0, 1\}^m \rightarrow \{0, 1\}$ and $S \subseteq [2^t]$ with $|S| \leq 2^k$.

Lemma 3.4.6 (Chain rule for min-entropy). *If $|\text{supp}(A)| \leq w$, then*

$$\tilde{H}_{\min}(X | A) \geq H_{\min}(X) - \log w.$$

If X has a lot of min-entropy given A , then intuitively, we should expect that $\text{Ext}(X, U_d)$ looks uniform even given A . This intuition is correct, as expressed by the following lemma, by Vadhan [Vad12]. (See also a similar earlier lemma by Dodis et al. [DORS08, Lemma 2.3].)

Lemma 3.4.7 (Extracting from sources with high conditional min-entropy [Vad12, Problem 6.8]). *Let $\text{Ext}: \{0, 1\}^t \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ be a (k, ε) -extractor. If $\tilde{H}_{\min}(X | A) \geq k$, then*

$$d_{\text{TV}}\left(\left(\text{Ext}(X, U_d), A\right), \left(U_m, A\right)\right) \leq 3\varepsilon.$$

(Here U_d is independent of (X, A) and U_m is independent of A .)

Finally, we will need the following standard “data processing inequality,” which says that applying a function – even a randomized function – can only make two distributions closer.

Lemma 3.4.8 (Data processing inequality for total variation distance). *Let A and \tilde{A} be random variables over the same space. Let R be independent of both A and \tilde{A} , and let f be any function. Then*

$$d_{\text{TV}}\left(f(A, R), f(\tilde{A}, R)\right) \leq d_{\text{TV}}\left(A, \tilde{A}\right).$$

3.4.2 Using extractors to fool standard-order ROBPs

We are now ready to analyze the PRG of Equation (3.1). We will show that it achieves the following parameters.

Lemma 3.4.9 (One iteration of the Nisan-Zuckerman generator). *Let $w, n \in \mathbb{N}$ with $n \geq \log w$, and let $\varepsilon > 0$. There is an explicit ε -PRG for width- w length- n standard-order ROBPs with seed length*

$$O\left(\log w + \frac{n \log(n/\varepsilon)}{\log w}\right).$$

Furthermore, the generator can be computed by an algorithm that reads the seed once from left to right and runs in space $O(\log(wn/\varepsilon))$.

Proof. Let $\ell = \frac{n}{\log w}$. Let $\varepsilon' = \frac{\varepsilon}{3\ell}$, and let $\text{Ext}: \{0, 1\}^{3\log w} \times \{0, 1\}^d \rightarrow \{0, 1\}^{\log w}$ be the $(2\log w, \varepsilon')$ -extractor of Theorem 3.4.4, so $d = O(\log(n/\varepsilon))$. The PRG G is given by Equation (3.1). The seed length and efficiency claims are clear.

As for correctness, our job is to show that G fools every width- w length- n standard-order ROBP. It will be convenient to group the n layers into ℓ blocks of $\log w$ layers each. This can be viewed as a width- w length- ℓ ROBP over the alphabet $\{0, 1\}^{\log w}$, i.e., each vertex has w outgoing edges labeled with all strings in $\{0, 1\}^{\log w}$. Let f be such a program, with layers V_0, \dots, V_ℓ . For $i < n$, $a \in V_i$, and $r \in \{0, 1\}^{\log w}$, let $a[r]$ denote the vertex reached from a by traversing the outgoing edge with label r .

Sample $X, Y_1, \dots, Y_\ell, R_1, \dots, R_\ell$ independently and uniformly at random, where $X \in \{0, 1\}^{3\log w}$, $Y_i \in \{0, 1\}^d$, and $R_i \in \{0, 1\}^{\log w}$. Let A_0, \dots, A_ℓ be the sequence of vertices reached when f reads the truly random bits R_1, \dots, R_ℓ , i.e., $A_0 = v_{\text{start}}$ and

$$A_i = A_{i-1}[R_i].$$

Similarly, let $\tilde{R}_i = \text{Ext}(X, Y_i)$, and let $\tilde{A}_0, \dots, \tilde{A}_\ell$ be the sequence of vertices reached when f reads the pseudorandom bits $\tilde{R}_1, \dots, \tilde{R}_\ell$, i.e., $\tilde{A}_0 = v_{\text{start}}$ and

$$\tilde{A}_i = \tilde{A}_{i-1}[\tilde{R}_i].$$

We will prove by induction on i that $d_{\text{TV}}(A_i, \tilde{A}_i) \leq 3\varepsilon' i$. The base case $i = 0$ is trivial. Now fix $i > 0$. By the triangle inequality, we have

$$\begin{aligned} d_{\text{TV}}(A_i, \tilde{A}_i) &= d_{\text{TV}}(A_{i-1}[\tilde{R}_i], \tilde{A}_{i-1}[\tilde{R}_i]) \\ &\leq d_{\text{TV}}(A_{i-1}[\tilde{R}_i], \tilde{A}_{i-1}[\tilde{R}_i]) + d_{\text{TV}}(\tilde{A}_{i-1}[\tilde{R}_i], \tilde{A}_{i-1}[\tilde{R}_i]). \end{aligned}$$

Applying the data processing inequality (Lemma 3.4.8) to each term, we get

$$d_{\text{TV}}(A_i, \tilde{A}_i) \leq d_{\text{TV}}(A_{i-1}, \tilde{A}_{i-1}) + d_{\text{TV}}((\tilde{A}_{i-1}, R_i), (\tilde{A}_{i-1}, \tilde{R}_i)).$$

The first term is at most $3\varepsilon' \cdot (i - 1)$ by induction. As for the second term, the chain rule for min-entropy (Lemma 3.4.6) implies that

$$\tilde{H}_{\min}(X \mid \tilde{A}_{i-1}) \geq 2 \log w.$$

Therefore, Lemma 3.4.7 guarantees that the second term is at most $3\varepsilon'$. Summing up, we get $d_{\text{TV}}(A_i, \tilde{A}_i) \leq 3\varepsilon' i$ as claimed. Therefore,

$$|\mathbb{E}[f] - \mathbb{E}[f(G(U_{3 \log w + \ell d}))]| = |\Pr[A_\ell \in V_{\text{accept}}] - \Pr[\tilde{A}_\ell \in V_{\text{accept}}]| \leq d_{\text{TV}}(A_\ell, \tilde{A}_\ell) \leq \varepsilon. \quad \square$$

Lemma 3.4.9 implies Theorem 3.4.1 when c is small, such as $c = 1.5$. To handle larger c , we compose the generator of Lemma 3.4.9 with itself. More details are below.

Proof of Theorem 3.4.1. Let $\varepsilon = 2^{-\log^{1-\nu} w}$. Let $w_1 = w$ and $n_1 = \log^c w$, and let G_1 be the ε -PRG of Lemma 3.4.9 for width- w_1 length- n_1 ROBPs. This generator G_1 has seed length $n_2 = O(\log^{c-\nu} w)$. Furthermore, G_1 can be computed by an algorithm that reads the seed once from left to right and runs in space $O(\log w)$. It follows that for any width- w standard-order ROBP f , the function $f_2(x) \stackrel{\text{def}}{=} f(G_1(x))$ can be computed by an ROBP of width $w_2 = \text{poly}(w)$. Therefore, let G_2 be the ε -PRG of Lemma 3.4.9 for width- w_2 length- n_2 standard-order ROBPs, with seed length $n_3 = O(\log^{c-2\nu} w)$. The composition $G_1 \circ G_2$ fools f with error 2ε . Once again, $f(G_1(G_2(x)))$ can be computed by a standard-order ROBP of width $w_3 = \text{poly}(w)$. Continuing in this way for $O(c/\nu) = O(1)$ steps, we obtain an explicit $O(\varepsilon)$ -PRG for f with seed length $O(\log w)$. \square

An interesting feature of the proof of Theorem 3.4.1 is that the *efficiency* of the PRG of Lemma 3.4.9 is a key part of the proof of *correctness* of the final PRG.

After Nisan and Zuckerman's work [NZ96], Armoni developed their techniques further and designed a PRG that “interpolates” between the INW generator and the Nisan-Zuckerman generator [Arm98]. Armoni's seed length was slightly improved later [KNW08] by plugging in extractors that were developed after Armoni's work [GUV09] (and analyzing their space complexity). At the extremes $n \geq w$ and $n \leq \text{polylog}(w)$, Armoni's generator has the same seed length as the INW generator and the Nisan-Zuckerman generator, respectively. However, in the intermediate regime $\text{polylog}(w) \ll n \ll w$, Armoni's generator (as optimized by Kane, Nelson, and Woodruff [KNW08]) is the best PRG known for ROBPs. It outperforms the INW generator in this regime by a factor of

$\log \log w$ (for moderate error), and this slight improvement (combined with other techniques) later led to the current best unconditional derandomization of space-bounded computation [Hoz21].

The Nisan-Zuckerman generator does not have optimal error. Improving the error is an appealing open problem.

Open Problem 3.4.10 (Improving the error of the Nisan-Zuckerman generator). *For some function $n = \omega(\log w)$, design an explicit PRG for width- w length- n standard-order ROBPs with seed length $O(\log w)$ and error $1/w$.*

The “hitting set generator” analogue of [Open Problem 3.4.10](#) has been solved [AKS87; HZ20].

Chapter 4

PRGs and Hardness

From an algorithm-design perspective, an explicit PRG construction is a “positive” theorem – an *upper bound* on the resources needed to sample a distribution that fools such-and-such class. An explicit PRG can be used as a building block in a larger algorithm.

On the other hand, from a complexity perspective, a construction of an explicit PRG G fooling a model \mathcal{F} provides a concrete example of a task that \mathcal{F} *cannot* do. A PRG construction is thus a “negative” theorem – an impossibility result – a *lower bound* on the resources required to distinguish the generator’s output from the uniform distribution.

In this chapter, we will explore connections between PRGs and more traditional “lower bound” notions. First, in [Section 4.1](#), we will investigate the “PRGs as lower bounds” viewpoint in more detail and discuss its implications for the prospect of PRG design. Then, in [Section 4.2](#), we will go the other direction, i.e., we will show how to construct PRGs from suitable lower bounds.

4.1 PRGs as high-quality lower bounds

4.1.1 PRGs imply hard Boolean functions

We begin by showing that for every PRG G for a class \mathcal{F} , there is a certain Boolean function h (closely related to G) that cannot be computed by a certain class \mathcal{F}' (closely related to \mathcal{F}). We will make the mild assumption that \mathcal{F} is “closed under restrictions” as defined below.

Definition 4.1.1 (Closure under restrictions). *Let \mathcal{F} be a class of functions f on $\{0,1\}^n$. We say that \mathcal{F} is closed under restrictions if the following holds. Let $f \in \mathcal{F}$, let $i \in [n]$, and let $b \in \{0,1\}$. Define $g(x) = f(x^{i \rightarrow b})$, where $x^{(i \rightarrow b)}$ denotes the string obtained from x by replacing the i -th bit with b . Then $g \in \mathcal{F}$.*

Proposition 4.1.2 (PRG \implies Hard Function). *Let \mathcal{F} be a class of functions $f: \{0,1\}^n \rightarrow \{0,1\}$ that is closed under restrictions. Let $G: \{0,1\}^s \rightarrow \{0,1\}^n$ be an ε -PRG for \mathcal{F} where $s < n$ and $\varepsilon < 1/2$. Define $h: \{0,1\}^{s+1} \rightarrow \{0,1\}$ by*

$$h(x) = 1 \iff \text{there exist } y, z \text{ such that } G(y) = (x, z).$$

Let \mathcal{F}' be the class of functions $f: \{0,1\}^{s+1} \rightarrow \{0,1\}$ of the form $f(x) = f_0(x, a)$ where $f_0 \in \mathcal{F}$ and $a \in \{0,1\}^{n-s-1}$. Then $h \notin \mathcal{F}'$.

Proof. Since G fools \mathcal{F} and \mathcal{F} is closed under restrictions, the distribution $G(U_s)_{1\dots s+1}$ fools \mathcal{F}' with error ε . In contrast, let us show that it does not ε -fool h . Indeed, $\mathbb{E}[h(G(U_s)_{1\dots s+1})] = 1$, whereas $\mathbb{E}[h] \leq 1/2$, because the image of G has size at most 2^s . \square

If the PRG G is explicit, then the hard function h in [Proposition 4.1.2](#) is “somewhat explicit.” For example, if $s = n^{\Omega(1)}$, then [Definition 1.4.1](#) implies that $h \in \mathbf{NP}$.

The conclusion of [Proposition 4.1.2](#) is a “worst-case” lower bound, i.e., it merely asserts that there is no function in \mathcal{F} that correctly computes the hard function g on *all* inputs. An “average-case” lower bound, as formalized below, is stronger: it asserts that no function in \mathcal{F} can correctly compute the hard function on significantly more than *half* of the inputs (with respect to some distribution).

Definition 4.1.3 (Average-case hardness). *Let \mathcal{F} be a class of functions $f: \{0, 1\}^r \rightarrow \{0, 1\}$, let $h: \{0, 1\}^r \rightarrow \{0, 1\}$, and let D be a distribution over $\{0, 1\}^r$. We say that h is ε -hard for \mathcal{F} with respect to D if for every $f \in \mathcal{F}$,*

$$\left| \Pr_{X \sim D} [f(X) = h(X)] - \frac{1}{2} \right| \leq \varepsilon.$$

(Note that in [Definition 4.1.3](#), we assume that the success probability is neither significantly more than $1/2$, nor significantly *less* than $1/2$. This is just for convenience. The two bounds are equivalent if \mathcal{F} is closed under complement.) The following refined reduction, first formalized by Viola [[Vio09](#)], shows that PRGs imply average-case hardness.

Proposition 4.1.4 (PRG \implies Average-Case Hardness [[Vio09](#)]). *Let \mathcal{F} be a class of functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$ that is closed under restrictions. Let $G: \{0, 1\}^s \rightarrow \{0, 1\}^n$ be an ε -PRG for \mathcal{F} . Let $r = s + \lceil \log(1/\varepsilon) \rceil$ and assume that $r \leq n$. Define $h: \{0, 1\}^r \rightarrow \{0, 1\}$ by*

$$h(x) = 1 \iff \text{there exist } y, z \text{ such that } G(y) = (x, z).$$

Let \mathcal{F}' be the class of functions $f: \{0, 1\}^r \rightarrow \{0, 1\}$ of the form $f(x) = f_0(x, a)$ where $f_0 \in \mathcal{F}$ and $a \in \{0, 1\}^{n-r}$. Let $D = \frac{1}{2}U_r + \frac{1}{2}G(U_s)_{1\dots r}$, i.e., the distribution D is a balanced convex combination of the distributions U_r and $G(U_s)_{1\dots r}$. Then h is ε -hard for \mathcal{F}' with respect to D .

Proof. Sample $U \sim U_r$ and $U' \sim U_s$. Fix any $f \in \mathcal{F}'$, say $f(x) = f_0(x, a)$. Then

$$\begin{aligned} \Pr_{X \sim D} [f(X) = h(X)] &= \frac{1}{2} \Pr[f(U) = h(U)] + \frac{1}{2} \Pr[f(G(U')_{1\dots r}) = h(G(U')_{1\dots r})] \\ &= \frac{1}{2} \Pr[f(U) = h(U)] + \frac{1}{2} \Pr[f(G(U')_{1\dots r}) = 1] \\ &\leq \frac{1}{2} \Pr[f(U) = 0] + \frac{1}{2} \Pr[h(U) = 1] + \frac{1}{2} (\Pr[f(U) = 1] + \varepsilon) \\ &= \frac{1}{2} + \frac{\varepsilon + \mathbb{E}[h]}{2} \\ &\leq \frac{1}{2} + \varepsilon, \end{aligned}$$

where the first inequality uses the fact that \mathcal{F} is closed under restrictions and hence the distribution $G(U')_{1\dots r}$ fools \mathcal{F}' with error ε . Since G also fools complements of functions in \mathcal{F} , we also get the reverse inequality $\Pr_{X \sim D} [f(X) = h(X)] \geq \frac{1}{2} - \varepsilon$. \square

[Proposition 4.1.4](#) demonstrates that there is a *hierarchy of lower bounds*:

$$\text{PRG} \implies \text{Average-Case Lower Bound} \implies \text{Worst-Case Lower Bound}. \quad (4.1)$$

A PRG construction is a lower bound that is particularly strong, qualitatively speaking. Admittedly, the average-case lower bound in [Proposition 4.1.4](#) is with respect to a certain (explicitly-sampleable) *non-uniform* distribution D , whereas traditionally, we seek average-case lower bounds with respect to the uniform distribution. However, it turns out that in many cases the two types of average-case lower bound are essentially equivalent [[CLLO21](#)].

4.1.2 The lack-of-lower-bounds barrier

Looking at Equation (4.1), an optimist might hope to use PRGs to prove new lower bounds. In practice, however, *lower bounds come first*. Therefore, a *lack* of known lower bounds for a particular model can be considered a type of *barrier* to constructing PRGs for that model.

For example, it would be great to obtain an explicit PRG that fools all size- n Boolean circuits with seed length, say, $O(\sqrt{n})$. By Proposition 4.1.2, such a PRG would imply the existence of a function $h \in \mathbf{NP}$ with quadratic circuit complexity. But proving that such a function exists would be a major breakthrough in circuit complexity. The conventional wisdom is that one should not try to design an unconditional PRG for size- n Boolean circuits until *after* proving the corresponding circuit lower bounds.

The good news, as we have seen already in Chapters 2 and 3, is that this “lack-of-lower-bounds barrier” still leaves plenty of room for a rich theory of unconditional PRGs. After all, highly nontrivial lower bounds are *already known* for many interesting classes, and hence we can try to design PRGs with matching parameters. For example, for \mathbf{AC}^0 circuits (see Definition 2.5.4), the state-of-the-art lower bounds are as follows.

Theorem 4.1.5 (Parity is hard for \mathbf{AC}^0 circuits [IMP12; Hås14]). *For every $m, d \in \mathbb{N}$ and $\varepsilon > 0$, there exists a value $r = O(\log^{d-1} m \cdot \log(1/\varepsilon))$ such that the parity function on r bits is ε -hard for depth- d size- m \mathbf{AC}^0 circuits with respect to the uniform distribution.*

In light of Theorem 4.1.5, we may reasonably hope to design an explicit ε -PRG for depth- d size- m \mathbf{AC}^0 circuits with seed length as low as

$$O(\log^{d-1} m \cdot \log(1/\varepsilon)). \quad (4.2)$$

(Assume for simplicity that $d \geq 2$ and $m \geq n$, where n is the number of pseudorandom bits.) Indeed, recall that Braverman’s theorem (Section 2.6) implies a fairly similar seed length of $\log^{O(d)} m \cdot \log(1/\varepsilon)$. Furthermore, as we will discuss in Section 5.3, explicit PRGs for \mathbf{AC}^0 are known with better seed lengths, getting very close to the bound of Equation (4.2). The optimal seed length would be $O(\log(m/\varepsilon))$, independent of d , but we should probably not expect to go below $O(\log^{d-1} m \cdot \log(1/\varepsilon))$ until *after* improving the known lower bounds for \mathbf{AC}^0 (Theorem 4.1.5).

As another example, let us consider standard-order ROBPs. Here the situation is better, because optimal lower bounds are known:

Proposition 4.1.6 (Inner product is hard for standard-order ROBPs). *For each even positive integer $2r \in \mathbb{N}$, let $\mathbf{IP}_{2r} : \{0, 1\}^{2r} \rightarrow \{0, 1\}$ denote the function*

$$\mathbf{IP}_{2r}(x, y) = \sum_{i=1}^r x_i y_i \pmod{2}.$$

For every $w \in \mathbb{N}$ and $\varepsilon > 0$, there exists a value $r = O(\log(w/\varepsilon))$ such that \mathbf{IP}_{2r} is ε -hard for width- w standard-order ROBPs with respect to the uniform distribution.

Proof. For every width- w length- $(2r)$ standard-order ROBP f , there is a communication protocol in which Alice holds x , Bob holds y , they communicate $1 + \lceil \log w \rceil$ bits, and then they output $f(x, y)$. (Alice simulates the first half of f and sends the state to Bob; Bob simulates the second half of f and sends the output bit to Alice.) The proposition follows by plugging in classic lower bounds on the average-case communication complexity of \mathbf{IP}_{2r} (e.g., see Rao and Yehudayoff’s text [RY20, Theorem 5.6]). \square

Thus, for standard-order ROBPs, there is no lack-of-lower-bounds barrier, and it is perfectly reasonable to try to design optimal PRGs.

4.2 The Nisan-Wigderson framework

In the previous section, we saw that a PRG implies a hard decision problem. In this chapter, we will discuss a famous method due to Nisan and Wigderson [NW94] for going the other way: converting a hard decision problem into a PRG. Thus, we will establish a fairly tight connection between PRGs and lower bounds (although there are losses that are important in some cases, so the connection is not a perfect equivalence).

4.2.1 Constructing a PRG from a hard function

Let $h: \{0,1\}^r \rightarrow \{0,1\}$ be a candidate “hard function.” Let $s > r$, and let S_1, \dots, S_n be a family of r -subsets of $[s]$, i.e., $S_i \subseteq [s]$ and $|S_i| = r$. The Nisan-Wigderson generator $G: \{0,1\}^s \rightarrow \{0,1\}^n$ is given by

$$G(x) = (h(x_{S_1}), \dots, h(x_{S_n})), \quad (4.3)$$

where $x_S = x_{i_1}x_{i_2} \dots x_{i_r}$ when $S = \{i_1 < i_2 < \dots < i_r\}$. In words, the generator applies the hard function to n different substrings of the seed.

We will prove that this construction works under certain assumptions on h and S_1, \dots, S_n . Intuitively, to ensure that the output bits of G appear to be independent, we should apply the hard function h to inputs that are almost “unrelated.” We will achieve this property by requiring that the sets of indices S_1, \dots, S_n are “nearly” disjoint.

Definition 4.2.1 (Nearly disjoint sets). *We say that sets S_1, \dots, S_n are k -nearly disjoint¹ if $|S_i \cap S_j| \leq k$ for all distinct $i, j \in [n]$.*

Meanwhile, the function h should be hard to compute, even on average (see Definition 4.1.3). To fool a function f , we will assume that h is hard for compositions of f with arbitrary k -juntas (see Definition 2.4.14). Under these assumptions, the Nisan-Wigderson generator achieves the following parameters.

Theorem 4.2.2 (Nisan-Wigderson reduction). *Let $f: \{0,1\}^n \rightarrow \{0,1\}$. Suppose $h: \{0,1\}^r \rightarrow \{0,1\}$ is ε -hard for $f \circ \text{JUNTA}_{r,k}$ with respect to the uniform distribution, and suppose that S_1, \dots, S_n are k -nearly disjoint r -subsets of $[s]$ for some $s > r$. Then the Nisan-Wigderson generator G given by Equation (4.3) fools f with error $\varepsilon \cdot n$.*

4.2.2 Analysis: Unpredictability

The proof of Theorem 4.2.2 is based on the problem of *predicting* the next bit of a pseudorandom string after seeing the first few bits. A truly random string is completely unpredictable.

Definition 4.2.3 (Unpredictability). *Let X be a distribution over $\{0,1\}^n$, let $f: \{0,1\}^n \rightarrow \{0,1\}$, and let $\varepsilon > 0$. We say that X is ε -unpredictable for f if for every $i \in [n]$ and every $a \in \{0,1\}^{n-i+1}$, we have*

$$\left| \Pr[f(X_1, X_2, \dots, X_{i-1}, a) = X_i] - \frac{1}{2} \right| \leq \varepsilon.$$

Equivalently, X fools the test $x \mapsto f(x_1, x_2, \dots, x_{i-1}, a) \oplus x_i$ with error ε . We say that a generator $G: \{0,1\}^s \rightarrow \{0,1\}^n$ is ε -unpredictable for f if $G(U_s)$ is ε -unpredictable for f .

¹A family of nearly disjoint sets, all of the same size, is also known as a *design* or a *partial design* or a *partial Steiner system* or a *packing*.

In some of the early literature, something like [Definition 4.2.3](#) is actually taken to be the definition of a PRG [[Sha83](#); [BM84](#)]. As a first step toward proving [Theorem 4.2.2](#), let us show that the Nisan-Wigderson generator is unpredictable.

Lemma 4.2.4 (The NW PRG is unpredictable). *Under the assumptions of [Theorem 4.2.2](#), the Nisan-Wigderson generator G is ε -unpredictable for f .*

Proof. Fix $i \in [n]$ and $a \in \{0, 1\}^{n-i+1}$. Sample a seed $Y \in \{0, 1\}^s$ uniformly at random, and let $X = G(Y)$. Then

$$\begin{aligned} \left| \Pr[f(X_1, \dots, X_{i-1}, a) = X_i] - \frac{1}{2} \right| &= \left| \mathbb{E}_{Y_{[s] \setminus S_i}} \left[\Pr_{Y_{S_i}}[f(h(Y_{S_1}), \dots, h(Y_{S_{i-1}}), a) = h(Y_{S_i})] \right] - \frac{1}{2} \right| \\ &\leq \mathbb{E}_{Y_{[s] \setminus S_i}} \left[\left| \Pr_{Y_{S_i}}[f(h(Y_{S_1}), \dots, h(Y_{S_{i-1}}), a) = h(Y_{S_i})] - \frac{1}{2} \right| \right]. \end{aligned}$$

Consider any arbitrary fixing of $Y_{[s] \setminus S_i}$ and let $Z = Y_{S_i}$. For each $j < i$, since we fixed $Y_{[s] \setminus S_i}$ and $|S_i \cap S_j| \leq k$, there is some k -junta ϕ_j such that $h(Y_{S_j}) = \phi_j(Z)$. Therefore,

$$\begin{aligned} \left| \Pr_{Y_{S_i}}[f(h(Y_{S_1}), \dots, h(Y_{S_{i-1}}), a) = h(Y_{S_i})] - \frac{1}{2} \right| &= \left| \Pr_Z[f(\phi_1(Z), \dots, \phi_{i-1}(Z), a) = h(Z)] - \frac{1}{2} \right| \\ &\leq \varepsilon, \end{aligned}$$

because h is ε -hard for $f \circ \text{JUNTA}_{r,k}$ (note that each bit of a can trivially be computed by a 0-junta.) \square

To complete the proof of [Theorem 4.2.2](#), we will relate the ‘‘predictor’’ model to the standard ‘‘distinguisher’’ model. We will show that if a distribution is unpredictable for f , then it also fools f , with a factor of n loss in the error parameter. This lemma is attributed to Yao.

Lemma 4.2.5 (Unpredictable \implies Pseudorandom). *Let X be a distribution over $\{0, 1\}^n$ and let $f: \{0, 1\}^n \rightarrow \{0, 1\}$. If X is ε -unpredictable for f , then X fools f with error $\varepsilon \cdot n$.*

Proof. Let $R \sim U_n$ be independent of X . Define hybrid distributions D_0, D_1, \dots, D_n by

$$D_i = X_1 X_2 \dots X_i R_{i+1} R_{i+2} \dots R_n,$$

so D_i consists of i pseudorandom bits followed by $n - i$ truly random bits. By the triangle inequality,

$$|\mathbb{E}[f(X)] - \mathbb{E}[f]| = |\mathbb{E}[f(D_n)] - \mathbb{E}[f(D_0)]| \leq \sum_{i=1}^n |\mathbb{E}[f(D_i)] - \mathbb{E}[f(D_{i-1})]|.$$

For each $i \in [n]$, we have

$$\begin{aligned} &|\mathbb{E}[f(D_i)] - \mathbb{E}[f(D_{i-1})]| \\ &= \left| \mathbb{E}[f(D_{i-1}) \mid R_i = X_i] - \left(\frac{1}{2} \mathbb{E}[f(D_{i-1}) \mid R_i = X_i] + \frac{1}{2} \mathbb{E}[f(D_{i-1}) \mid R_i \neq X_i] \right) \right| \\ &= \left| \frac{1}{2} \mathbb{E}[f(D_{i-1}) \mid R_i = X_i] + \frac{1}{2} \mathbb{E}[\neg f(D_{i-1}) \mid R_i \neq X_i] - \frac{1}{2} \right| \\ &= \left| \mathbb{E}[f(D_{i-1}) \oplus R_i \oplus X_i] - \frac{1}{2} \right| \\ &\leq \mathbb{E}_R \left[\left| \mathbb{E}_X[f(D_{i-1}) \oplus R_i \oplus X_i] - \frac{1}{2} \right| \right]. \end{aligned}$$

This is at most ε , because for any fixing of R , if we let $g(x) = f(x_1 \dots x_{i-1} R_i \dots R_n) \oplus R_i \oplus x_i$, then either g or $\neg g$ is testing whether f successfully predicts x_i given x_1, \dots, x_{i-1} . Summing up, we get $|\mathbb{E}[f(X)] - \mathbb{E}[f]| \leq \varepsilon \cdot n$. \square

4.2.3 A family of nearly disjoint sets

We have now shown how to construct a PRG given two ingredients: a hard function h and a family of nearly disjoint sets S_1, \dots, S_n . The hard function must be tailored to the specific class of functions that we wish to fool, but constructing the family of nearly disjoint sets is a combinatorics problem that can be addressed separately. We will construct such a family using an argument by Erdős, Frankl, and Füredi [EFF85].

Lemma 4.2.6 (Existence of nearly disjoint sets [EFF85]). *Let $k, r, n \in \mathbb{N}$ with $r > k \geq 1$. For a suitable value $s = O(n^{1/(k+1)} \cdot r^2/k)$, there exists a k -nearly disjoint family of r -subsets $S_1, \dots, S_n \subseteq [s]$. Furthermore, given k, r , and n , the family can be deterministically constructed in time $\text{poly}(n, 2^s)$.*

Proof. Construct S_1, \dots, S_n greedily. That is, for $i \in [n]$, having constructed S_1, \dots, S_{i-1} , search exhaustively through all subsets of $[s]$ to find a set S_i of size r such that for every $j < i$, we have $|S_j \cap S_i| \leq k$. To prove that such a set exists, consider picking S_i uniformly at random from among all subsets of $[s]$ of size r . Then by the union bound,

$$\begin{aligned} \Pr_{S_i}[\exists j < i \text{ such that } |S_i \cap S_j| > k] &\leq \sum_{j=1}^{i-1} \Pr_{S_i}[|S_j \cap S_i| > k] = (i-1) \cdot \frac{\binom{r}{k+1} \cdot \binom{s-(k+1)}{r-(k+1)}}{\binom{s}{r}} \\ &< n \cdot \frac{\binom{r}{k+1}^2}{\binom{s}{k+1}}, \end{aligned}$$

where the last step uses the identity $\binom{s}{r} \cdot \binom{r}{k+1} = \binom{s}{k+1} \cdot \binom{s-(k+1)}{r-(k+1)}$. Therefore, a suitable S_i is guaranteed to exist provided $n \leq \binom{s}{k+1} / \binom{r}{k+1}^2$. Choose the universe size to be $s = \lceil er^2 \cdot n^{1/(k+1)} / k \rceil$, because that way

$$n \leq \left(\frac{sk}{er^2} \right)^{k+1} < \frac{\binom{s}{k+1}}{\binom{r}{k+1}^2},$$

where the last step uses $0 < k < r \leq s$. \square

Optimality Ignoring explicitness, what is the best possible universe size s in Lemma 4.2.6? Equivalently, given k, r , and s , what is the largest number n such that there exists a k -nearly disjoint family of r -subsets $S_1, \dots, S_n \subseteq [s]$? This seems to be an open problem in combinatorics, even if we are only interested in rough asymptotics. The proof of Lemma 4.2.6 shows the existence of a family with $n \geq \binom{s}{k+1} / \binom{r}{k+1}^2$. Conversely, in every such family, each size- $(k+1)$ subset of $[s]$ is contained in at most one S_i , so² $n \leq \binom{s}{k+1} / \binom{r}{k+1}$. This last bound can be slightly improved [Sch64; Für88], but it seems that there is still a significant gap.

Open Problem 4.2.7 (Optimal nearly disjoint sets). *For given values k, r , and s , determine (asymptotically) the maximum value n such that there exists a k -nearly disjoint family of r -subsets $S_1, \dots, S_n \subseteq [s]$.*

²In terms of universe size, we get $s \geq \Omega(n^{1/(k+1)} \cdot r)$. One can also prove $s \geq \Omega(\min\{r^2/k, nr\})$ using the inclusion-exclusion principle.

When k and r are constant, the problem has been solved: a famous theorem by Rödl [Röd85] says that there exist families with $n = (1 - o(1)) \cdot \binom{s}{k+1} / \binom{r}{k+1}$. When k and r are growing parameters, however, the situation seems to be less clear. For example, when $r = k^2$ and $s = 100k^3$, the optimal value of n is somewhere between $2^{\Theta(k)}$ and $k^{\Theta(k)}$, but the true value seems to be unknown.

Efficiency The proof of Lemma 4.2.6 is simple, but it’s somewhat unsatisfactory because of the exhaustive search. The universe size s corresponds to the seed length of the Nisan-Wigderson generator. The time complexity $\text{poly}(n, 2^s)$ in Lemma 4.2.6 is too high to get a strictly “explicit” PRG, except in the case $s = O(\log n)$, since our definition of explicitness (Definition 1.4.1) is that the runtime should be $\text{poly}(n)$. In the literature [NW94; KM02; HR03], there are several constructions of families of nearly disjoint sets that are “more explicit” than the construction of Lemma 4.2.6, but unfortunately, the parameters of these constructions are not quite as good.

Open Problem 4.2.8 (More efficient constructions of nearly disjoint sets). *Find a family of nearly disjoint sets that has the same parameters as Lemma 4.2.6 and that can be constructed in time $\text{poly}(n)$. (Assume $s < n$.)*

4.2.4 Unconditional applications

To illustrate the Nisan-Wigderson framework, let us use the framework to design another PRG for \mathbf{AC}^0 circuits. Recall that the parity function is hard for such circuits (Theorem 4.1.5). By plugging the parity function into the Nisan-Wigderson framework, we get a PRG with the following parameters.

Corollary 4.2.9 (Hardness-based PRG for \mathbf{AC}^0). *For any $n, m, d \in \mathbb{N}$ and $\varepsilon > 0$, there is an ε -PRG for depth- d size- m \mathbf{AC}^0 circuits on n input bits with seed length $s = \log^{2d+O(1)}(mn) \cdot \text{polylog}(1/\varepsilon)$. The PRG can be computed in time $2^{O(s)}$.*

Proof. Let $k = \lfloor \log n \rfloor$. By Theorem 4.1.5, for a suitable value

$$r = O(\log^{d+1}(mn) \cdot \log(n/\varepsilon)),$$

the parity function $h: \{0, 1\}^r \rightarrow \{0, 1\}$ is (ε/n) -hard for depth- $(d+2)$ size- $(m+n \cdot 2^k)$ \mathbf{AC}^0 circuits. Let $S_1, \dots, S_n \subseteq [s]$ be the k -nearly disjoint family of r -subsets from Lemma 4.2.6, and let G be the Nisan-Wigderson PRG given by Equation (4.3). To prove the correctness of this PRG, let f be a depth- d size- m \mathbf{AC}^0 circuit on n input bits. Every function in $f \circ \text{JUNTA}_{r,k}$ can be computed by a circuit of depth $d+2$ and size $m+n \cdot 2^k$, because every k -junta can be computed by a depth-two \mathbf{AC}^0 circuit of size 2^k . Therefore, h is (ε/n) -hard for $f \circ \text{JUNTA}_{r,k}$, so Theorem 4.2.2 implies that G fools f with error ε . By Lemma 4.2.6, the seed length of our generator is $s = O(n^{1/(k+1)} \cdot r^2/k)$, which is $\log^{2d+O(1)}(mn) \cdot \log^2(1/\varepsilon)$ as claimed. \square

Historically, the Nisan-Wigderson approach provided the first explicit PRG for constant-depth polynomial-size \mathbf{AC}^0 circuits with seed length $\text{polylog } n$ [Nis91]. The seed length in Corollary 4.2.9 is a little better than the seed length implied by Braverman’s theorem (Section 2.6) in some cases, because the factor 2 in the exponent is better. Today, we can use other frameworks to construct PRGs for \mathbf{AC}^0 circuits with better seed lengths (see the discussion in Section 5.3.) However, the Nisan-Wigderson framework remains a valuable, flexible approach for designing PRGs, especially for more powerful models of computation. For example, the Nisan-Wigderson framework has been used to construct unconditional PRGs for \mathbf{AC}^0 circuits augmented with a few gates that compute arbitrary threshold functions or symmetric functions [LVW93; Vio07; LS11; ST18; KL18]. Also, a line of work initiated by Trevisan [Tre01] shows that there are connections between the Nisan-Wigderson framework and unconditional constructions of *randomness extractors* (see Definition 3.4.3).

4.3 Hardness-based PRGs beyond Nisan-Wigderson

In summary, the Nisan-Wigderson framework is a method for converting a hard function into a PRG. Starting from a function on r bits that is ε -hard for $f \circ \text{JUNTA}_{r,k}$, we get a PRG for f with seed length $O(n^{1/(k+1)} \cdot r^2/k)$ and error $\varepsilon \cdot n$.

Many other PRG constructions, including those that we saw in [Chapters 2 and 3](#) and those that we will see in [Chapter 5](#), are related to lower bounds in a less direct way. To design a PRG for a class \mathcal{F} , rather than using the mere fact that such-and-such lower bound holds, we can try to distill and develop the *insights* that were used to *prove* the lower bound. For example, as we saw in [Section 2.6](#), Braverman’s theorem relies on the LMN theorem ([Theorem 2.6.2](#)), which builds on Håstad’s switching lemma [[Has89](#)], which was originally proven for the sake of proving lower bounds for AC^0 [[Has89](#)]. As another example, the INW generator for ROBPs ([Section 3.2](#)) relies on the same “communication complexity” intuition that appears in the proof of optimal lower bounds for ROBPs ([Proposition 4.1.6](#)).

There are also methods other than the Nisan-Wigderson framework for generically converting hardness into randomness. These methods improve on the Nisan-Wigderson framework in fascinating and important ways. However, the known applications of these other methods are almost exclusively *conditional*. Since our focus is on unconditional PRGs, we will only briefly survey these other methods.

Optimizing the circuit-size blow-up Suppose we wish to design a PRG for size- n circuits of unbounded depth. The Nisan-Wigderson framework can produce such a PRG, given a function that is hard for circuits that are a little larger. Indeed, if f is a size- n circuit, then every function in $f \circ \text{JUNTA}_{r,k}$ can be computed by a circuit of size $m = n \cdot 2^{O(k)}$. To avoid paying the severe $n^{1/(k+1)}$ penalty in the Nisan-Wigderson seed length, one can choose $k = \Theta(\log n)$, in which case $m = n^{1+\Theta(1)}$.

There is a line of work on improving the size complexity m , i.e., showing how to construct a PRG for size- n circuits given a function that is hard for (some type of) circuits of size $O(n)$ or even $(1 + \alpha)n$ [[RRV02](#); [HR03](#); [DMOZ20](#); [CT21](#)].

Read-once models There is another “loss” in the Nisan-Wigderson reduction of a more qualitative nature. The Nisan-Wigderson framework is not well-suited for the important problem of fooling standard-order ROBPs. The reason is that every polynomial-size *read-many* branching program can be written in the form $f \circ \text{JUNTA}_{r,1}$ where f is a polynomial-size standard-order ROBP. Read-many branching programs are vastly more powerful than ROBPs, and the Nisan-Wigderson framework does not give us any way to take advantage of the read-once condition.

Babai, Nisan, and Szegedy nevertheless designed an unconditional hardness-based PRG for standard-order ROBPs [[BNS92](#)] via a different reduction. (Their generator was superseded by superior PRGs such as Nisan’s PRG [[Nis92](#)] and the INW PRG [[INW94](#)], which we discussed in [Section 3.2](#).)

The seed length compared to the domain size The seed length in the Nisan-Wigderson reduction is not ideal. Recall that a PRG with seed length s implies a hard function on $s + 1$ bits ([Proposition 4.1.2](#)). Therefore, starting from a hard function on r bits, we can hope to construct a PRG with seed length roughly r , rather than the r^2 factor that appears in [Lemma 4.2.6](#). Indeed, there is a line of work showing how to convert an appropriately-hard function on r bits into a PRG with seed length $O(r)$ or even $(1 + \alpha)r$ [[ISW06](#); [SU05](#); [Uma03](#); [DMOZ20](#); [CT21](#)].

Worst-case hardness assumptions Recall that in the Nisan-Wigderson framework, we rely on access to a function h that is hard on average, i.e., it is hard to compute h on even a $(1/2 + \varepsilon)$ -fraction of inputs. There is a line of work on constructing PRGs from *worst-case* hardness [IW97; STV01; ISW06; SU05; Uma03; DMOZ20; CT21]. To highlight one example, Umans showed how to construct a PRG for size- n circuits given a function that cannot be computed (in the worst case) by circuits of size n^c for a suitable constant c [Uma03]. If the hard function is on r bits, then the PRG has seed length $O(r)$, and given the truth table of the hard function and a seed, the PRG can be computed in time $2^{O(r)}$ [Uma03].

Uniform hardness assumptions In the Nisan-Wigderson framework, we start with a function h on a finite domain that is hard for some concrete, “nonuniform” model of computation. There are also many known constructions of PRGs from *uniform* complexity-theoretic hardness assumptions such as $\mathbf{BPP} \neq \mathbf{EXP}$ [BFNW93; CNS99; IW01; TV07; Gol11; CIS18; CRTY20].

The error parameter The Nisan-Wigderson reduction converts an ε -hard function into a PRG with error $\varepsilon \cdot n$. Fefferman, Shaltiel, Umans, and Viola have studied the problem of avoiding the factor-of- n blow-up [FSUV13].

One of the motivations for studying this problem is the challenge of designing better PRGs for $\mathbf{AC}^0[\oplus]$ circuits, i.e., \mathbf{AC}^0 circuits augmented with parity gates. For context, fairly strong lower bounds are known for this class. In particular, it follows from Smolensky’s work [Smo93; Fil10] that for every $m, d \in \mathbb{N}$ and $\varepsilon > 0$ with $d \leq \frac{\log(m/\varepsilon)}{\log \log(m/\varepsilon)}$, there is a value $r = \varepsilon^{-2} \cdot (\log(m/\varepsilon))^{O(d)}$ such that the majority function on r bits is ε -hard for depth- d size- m $\mathbf{AC}^0[\oplus]$ circuits. Therefore, even if we treat the lack-of-lower-bounds barrier as a real barrier, we can hope to design PRGs for these circuits with polylogarithmic seed length (at least in the constant-error regime). So far, however, the best fully-explicit PRGs for these circuits have much larger seed length, very close to the trivial seed length of n bits [FSUV13].

Open Problem 4.3.1 (PRGs for $\mathbf{AC}^0[\oplus]$). *Design an explicit PRG that fools constant-depth polynomial-size $\mathbf{AC}^0[\oplus]$ circuits on n input variables with seed length $o(n)$.*

Note that if we severely relax our standards of “explicitness,” then there are known constructions of PRGs for $\mathbf{AC}^0[\oplus]$ with seed length $n^{o(1)}$ [CR22; CLW20].

Cryptographic PRGs The Nisan-Wigderson generator is unsuitable for cryptography, because computing the generator involves evaluating the “hard function” h . A cryptographic PRG cannot afford to evaluate a function that is hard for the adversary to compute, because a cryptographic PRG must fool all efficient adversaries – including those that use polynomially more time than the PRG itself uses.

Nevertheless, the paradigm of using some type of “hard function” to construct a PRG has been highly successful in the cryptographic setting [Yao82; GMT82; Sha83; BM84; BBS86; Lev87; ACGS88; GL89; GKL93; HILL99; Gol01; Hol06; HHR06; HHR11; VZ12; HRV13; YLW15; YGLW15; DVV16; EWT21; MZ21]. Indeed, hardness-based cryptographic PRGs predate the Nisan-Wigderson framework.

Chapter 5

Random Restrictions

In this chapter, we will present several constructions of PRGs based on *random restrictions*. A *restriction* is a string $R \in \{0, 1, \star\}^n$. Intuitively, when $R_i = \star$, the interpretation is that R does not assign any value to the i -th variable. A restriction R can be *applied* to a function f via the following two definitions.

Definition 5.0.1 (Composition of restrictions). *If $R \in \{0, 1, \star\}^n$ and $x \in \{0, 1\}^n$, define the composition $R \circ x \in \{0, 1, \star\}^n$ by*

$$(R \circ x)_i = \begin{cases} R_i & \text{if } R_i \in \{0, 1\} \\ x_i & \text{if } R_i = \star. \end{cases} \quad (5.1)$$

More generally, the same formula holds when x is an element of $\{0, 1, \star\}^n$ rather than $\{0, 1\}^n$, so \circ is an associative binary operation on the space $\{0, 1, \star\}^n$.

Definition 5.0.2 (Applying a restriction to a function). *Let f be a function on $\{0, 1\}^n$, and let $R \in \{0, 1, \star\}^n$ be a restriction. Then the restricted function $f|_R$ is the function on $\{0, 1\}^n$ defined by*

$$f|_R(x) = f(R \circ x).$$

Remark 5.0.3 (Order of restriction composition). *Some sources define the composition operator \circ the other way around, so $f|_R(x) = f(x \circ R)$ rather than $f(R \circ x)$. Both conventions are reasonable. The motivation behind our convention is that one can identify a restriction $R \in \{0, 1, \star\}^n$ with the unique function $R: \{0, 1\}^n \rightarrow \{0, 1\}^n$ such that for every f and x , we have $f|_R(x) = f(R(x))$. (See [Figure 5.1](#).) Under this identification, the restriction composition operator \circ is literally function composition.*

We will often consider *random* restrictions as defined below.

Definition 5.0.4 (Truly random restrictions). *For $p > 0$, let \mathcal{R}_p denote a truly random restriction over $\{0, 1, \star\}^n$ with \star -probability p , i.e., the coordinates are independent, and each coordinate is*

$$\begin{cases} \star & \text{with probability } p \\ 0 & \text{with probability } (1-p)/2 \\ 1 & \text{with probability } (1-p)/2. \end{cases}$$

(The parameter n will be clear from context.)

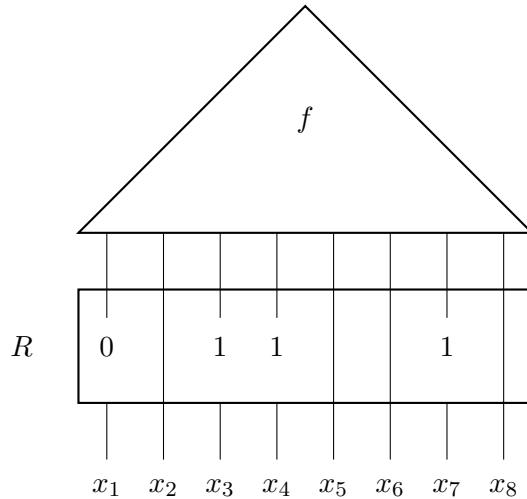


Figure 5.1: The restricted function $f|_R$ in the case $R = \{0, \star, 1, 1, \star, \star, 1, \star\}$. We still think of $f|_R$ as a function on the 8-variable domain $\{0, 1\}^8$, but its output only depends on the values of four of those variables.

Random restrictions have been used in many areas of the theory of computing, perhaps starting with Subbotovskaya’s pioneering work on De Morgan formulas [Sub61]. The process of designing a PRG for a class \mathcal{F} using restrictions can be divided into two main steps.

1. Prove a lemma that says that functions in \mathcal{F} *simplify* in some sense under restrictions. We will refer to such a lemma as a “simplification-under-restrictions lemma.” This first step requires an intimate understanding of the specific class \mathcal{F} .
2. Apply a generic *reduction* that says how to construct a PRG for any class satisfying a suitable simplification-under-restrictions lemma (and perhaps also satisfying some mild conditions such as closure properties). This second step is mainly about the abstract logic of restrictions and PRGs.

As we will see, this plan can be instantiated in multiple ways. There are *several different types* of simplification-under-restrictions lemmas:

- Are we considering truly random restrictions, or pseudorandom restrictions, or “partially pseudorandom” restrictions?
- What is our measure of “simplicity?”
- Does simplification occur with high probability, or does it merely occur “on average?”

Correspondingly, there are *several distinct reductions* from the problem of constructing PRGs to the problem of proving simplification under restrictions. We will discuss the *polarizing random walks* framework (Sections 5.1 and 5.2), the *iterated restrictions* paradigm (Sections 5.3 to 5.5), and the *Impagliazzo-Meka-Zuckerman* framework (Section 5.6). Perhaps these variations make the topic a bit confusing, but on the bright side, all this flexibility means that we have a rich toolkit for constructing PRGs.

5.1 PRGs from polarizing random walks

In this section, we present our first reduction showing that if a class simplifies under random restrictions, then we get a PRG for that class. This first reduction, based on “polarizing random walks,” was introduced relatively recently by Chattopadhyay, Hatami, Hosseini, and Lovett [CHHL19]. It has the benefit that its assumptions are quite minimal, i.e., it is applicable in a relatively broad set of circumstances.

5.1.1 Simplification under truly random restrictions

Let \mathcal{F} be a class of functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$ that we wish to fool. Suppose that we have shown that functions in our class \mathcal{F} simplify under random restrictions. Specifically, suppose we have identified a class $\mathcal{F}_{\text{simp}}$ of “simpler” functions and values $p, \delta > 0$ such that for each $f \in \mathcal{F}$, we have

$$\Pr[f|_{\mathcal{R}_p} \in \mathcal{F}_{\text{simp}}] \geq 1 - \delta.$$

For example, Rossman proved the following theorem [Ros19] using Håstad’s famous switching lemma [Has89] and his more recent “multi-switching” lemma [Hås14].

Theorem 5.1.1 (\mathbf{AC}^0 simplifies under restrictions [Ros19]). *For every $n, m, d \in \mathbb{N}$, there is a value $p = 1/\Theta(\log m)^{d-1}$ such that if $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is computable by a depth- d size- m \mathbf{AC}^0 circuit and $\delta > 0$, then with probability $1 - \delta$, the restricted function $f|_{\mathcal{R}_p}$ can be computed by a decision tree of depth at most $\log(1/\delta)$.*

For this section, we will consider a class “simple” if we can fool it with a short seed. For example, we consider small-depth decision trees to be simple, because we can ε -fool decision trees of depth $\log(1/\delta)$ using a seed of length $O(\log(1/\delta) + \log(1/\varepsilon) + \log \log n)$ (see Section 2.3.3). In general, assuming that functions in \mathcal{F} simplify to $\mathcal{F}_{\text{simp}}$ under restrictions, and assuming that we have a good PRG for $\mathcal{F}_{\text{simp}}$, how can we design a PRG for the original class \mathcal{F} ? Our strategy will be to first design a relaxation of a PRG called a *fractional PRG*. Then, we will gradually transform the fractional PRG into a genuine PRG.

5.1.2 Fractional PRGs

For the purposes of this approach, it will be convenient to work with Boolean functions $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ instead of our usual domain $\{0, 1\}^n$. Recall that a PRG for a family of Boolean functions is a function $G: \{0, 1\}^s \rightarrow \{-1, 1\}^n$ such that $\mathbb{E}[f(G(U_s))] \approx \mathbb{E}[f]$ for every function f in the family. A *fractional PRG* is a relaxation of a PRG where we allow G to take values in the solid hypercube $[-1, 1]^n$ as opposed to $\{-1, 1\}^n$. For this to make sense, we would like to be able to “evaluate” $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ on arbitrary inputs from $[-1, 1]^n$. One natural way to do this is by considering the multilinear extension of f . Recall that every Boolean function $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ has a unique Fourier expansion,

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{i \in S} x_i.$$

The formula above can be evaluated at an arbitrary point $x \in \mathbb{R}^n$, which allows us to extend f to a multilinear polynomial $f: \mathbb{R}^n \rightarrow \mathbb{R}$. To understand how to interpret the value $f(x)$ for fractional vectors $x \in [-1, 1]^n$, we make the following definition.

Definition 5.1.2 (Product distribution notation). *For $x \in [-1, 1]^n$, let Π_x be the unique product distribution over $\{-1, 1\}^n$ satisfying $\mathbb{E}[\Pi_x] = x$.*

Fact 5.1.3 (Evaluation at fractional points). *Let $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ be any function. Extend f to the domain $[-1, 1]^n$ via the Fourier expansion. Then for every $x \in [-1, 1]^n$, we have*

$$f(x) = \mathbb{E}[f(\Pi_x)].$$

More generally, for every product distribution X over $[-1, 1]^n$, we have $\mathbb{E}[f(X)] = f(\mathbb{E}[X])$.

Proof. This is immediate from multilinearity and linearity of expectation. \square

As a result, if f takes values in $\{-1, 1\}$, then its multilinear extension is a map $[-1, 1]^n \rightarrow [-1, 1]$. Another useful corollary is that $f(0^n) = \mathbb{E}[f]$. Thus, a PRG can be seen as a method of sampling a distribution X over $\{-1, 1\}^n$ such that $\mathbb{E}[f(X)] \approx f(0^n)$. We will define a fractional PRG by allowing the pseudorandom distribution to take values inside the cube.

Definition 5.1.4 (Fractional PRGs). *Let $f: \{-1, 1\}^n \rightarrow \mathbb{R}$, and extend f to the domain $[-1, 1]^n$ via the Fourier expansion. A random variable $X \in [-1, 1]^n$ is said to ε -fool f , or fool f with error ε , if*

$$|\mathbb{E}[f(X)] - f(0^n)| \leq \varepsilon.$$

We say that X fools a family \mathcal{F} of Boolean functions with error ε if X fools (the multilinear extension of) each function in \mathcal{F} with error ε . A fractional ε -PRG for \mathcal{F} is a function $G: \{0, 1\}^s \rightarrow [-1, 1]^n$ such that $G(U_s)$ fools \mathcal{F} with error ε .

We can trivially fool *all* Boolean functions with error 0 and seed length 0 by simply outputting 0^n . However, our main motivation for defining fractional PRGs is as a means to constructing true PRGs, and our PRG construction will require some non-triviality conditions from the fractional PRG. In particular, we will require each coordinate of the fractional PRG to have variance bounded away from zero.

Definition 5.1.5 (Noticeability). *We say that a random variable $X \in [-1, 1]^n$ is q -noticeable for a parameter $q \geq 0$, if for every $i \in [n]$, $\mathbb{E}[X_i^2] \geq q$. We say that a fractional PRG $G: \{0, 1\}^s \rightarrow [-1, 1]^n$ is q -noticeable if $G(U_s)$ is q -noticeable.*

The following lemma shows that if a class \mathcal{F} simplifies to another class $\mathcal{F}_{\text{simp}}$ under restrictions, and we have a good PRG for $\mathcal{F}_{\text{simp}}$, then we get a good *fractional* PRG for the original class \mathcal{F} , where the noticeability depends on the \star -probability of the restrictions.

Lemma 5.1.6 (Simplification implies fractional PRGs). *Let \mathcal{F} and $\mathcal{F}_{\text{simp}}$ be classes of functions $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$. Let $p, \delta > 0$, and suppose that for each $f \in \mathcal{F}$, we have*

$$\Pr[f|_{\mathcal{R}_p} \in \mathcal{F}_{\text{simp}}] \geq 1 - \delta.$$

Let X be a distribution over $\{-1, 1\}^n$ that ε -fools $\mathcal{F}_{\text{simp}}$. Then pX is (p^2) -noticeable and fools \mathcal{F} with error $\varepsilon + 2\delta$.

Proof. Clearly, we always have $(pX)_i^2 = p^2$, showing that pX is (p^2) -noticeable. Now fix $f \in \mathcal{F}$, and sample $R \sim \mathcal{R}_p$ independently of X . For each fixed string $x \in \{\pm 1\}^n$, the composition $R \circ x$ is a product distribution over $\{\pm 1\}^n$, where

$$\mathbb{E}[(R \circ x)_i] = (1 - p) \cdot 0 + p \cdot x_i = p \cdot x_i.$$

Therefore, by [Fact 5.1.3](#), we have $\mathbb{E}[f(R \circ x)] = f(px)$. Consequently, $\mathbb{E}_X[f(pX)] = \mathbb{E}_{R, X}[f(R \circ X)]$. Clearly, $|\mathbb{E}_{R, X}[f(R \circ X)] - \mathbb{E}[f]| \leq \varepsilon + 2\delta$. \square

By combining [Lemma 5.1.6](#) and [Theorem 5.1.1](#), we get a fractional PRG for \mathbf{AC}^0 with the following parameters.

Corollary 5.1.7 (Fractional PRGs for \mathbf{AC}^0). *For every $n, m, d \in \mathbb{N}$ and every $\varepsilon > 0$, there is an explicit q -noticeable fractional PRG that ε -fools depth- d size- m \mathbf{AC}^0 circuits with seed length $O(\log(1/\varepsilon) + \log \log n)$, where $q = 1/\Theta(\log m)^{2d-2}$.*

5.1.3 From fractional PRGs to PRGs

In this section we will prove that fractional PRGs can be used to construct PRGs with small seed length as long as the fractional PRG has two useful properties: it has a small seed length and all its coordinates have noticeable variance.

Theorem 5.1.8 (Fractional PRG \implies Standard PRG [[CHHL19](#)]). *Suppose that \mathcal{F} is a family of functions $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ that is closed under restrictions and shifts. Assume that there exists an explicit q -noticeable fractional PRG for \mathcal{F} with error ε and seed length s . Then there exists an explicit PRG for \mathcal{F} with seed length $O(s \cdot q^{-1} \cdot \log(n/\varepsilon))$ and error $O(\varepsilon \cdot q^{-1} \cdot \log(n/\varepsilon))$.*

For example, by combining [Theorem 5.1.8](#) and [Corollary 5.1.7](#), we get the following PRG for \mathbf{AC}^0 .

Corollary 5.1.9 (PRG for \mathbf{AC}^0 based on simplification under truly random restrictions). *For every $n, m, d \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for depth- d size- m \mathbf{AC}^0 circuits on n input bits with seed length*

$$\tilde{O}(\log m)^{2d-2} \cdot \tilde{O}(\log(n/\varepsilon) \cdot \log(1/\varepsilon)).$$

The seed length in [Corollary 5.1.9](#) is slightly better than the other PRGs for \mathbf{AC}^0 circuits that we have already seen (Braverman’s theorem in [Section 2.6](#), and the Nisan-Wigderson generator in [Section 4.2](#)). More importantly, this new PRG generalizes in some ways that the previously discussed PRGs do not. After all, the same approach works whenever a class simplifies under random restrictions. In fact, as we will discuss in [Section 5.2](#), it even works in the more general setting of a class that simplifies “on average” under random restrictions.

Random walks

To prove [Theorem 5.1.8](#), we will take a *random walk* through the solid hypercube $[-1, 1]^n$. It is natural to take $Y^{(0)} = 0^n$ as the starting point, since $\mathbb{E}[f] = f(0^n)$. Our goal then is to define a random walk that converges quickly to the Boolean cube $\{-1, 1\}^n$, while each step of the walk does not incur much error. We will define the steps by independent samples from the output distribution X of the fractional PRG.

To this end let $X^{(1)}, \dots, X^{(t)}$ be t independent samples of X where t is to be determined later. A natural first step for the random walk is $Y^{(1)} = Y^{(0)} + X^{(1)}$, as it has the two useful properties:

1. $|\mathbb{E}[f(Y^{(1)})] - \mathbb{E}[f(Y^{(0)})]| \leq \varepsilon$, and
2. Each coordinate of $Y^{(1)}$ is likely closer to $\{-1, 1\}$ due to p -noticeability.

It is tempting to continue this approach for all steps and in particular define $Y^{(j)} = Y^{(j-1)} + X^{(j)}$. This does not work, since we may already step out of the $[-1, 1]^n$ cube, and in fact after some steps start getting farther and farther from $\{-1, 1\}^n$. A slight modification that works is to normalize coordinates according to their distance from $\{-1, 1\}^n$.

For two vectors $x, x' \in [-1, 1]^n$, define $x \odot x' \in [-1, 1]^n$ to be their coordinate-wise product. Moreover, for every vector $y \in [-1, 1]^n$ define $\delta_y \in [0, 1]^n$ to be the vector with i -th coordinate $(\delta_y)_i = 1 - |y_i|$, i.e., $(\delta_y)_i$ is the distance from $y_i \in [-1, 1]$ to the Boolean endpoints $\{-1, 1\}$. The vector δ_y defines dimensions of the largest subcube inside $[-1, 1]^n$ centered at y . Using this notation, we can now define the random walk:

- $Y^{(0)} = 0^n$, and
- For $j > 0$, let $Y^{(j)} = Y^{(j-1)} + \delta_{Y^{(j-1)}} \odot X^{(j)}$.

We will show that this random walk quickly gets close to $\{-1, 1\}^n$. Still, there is a chance that the coordinates of $Y^{(t)}$ are never *exactly* integers. The final construction takes care of this by outputting the coordinate-wise signs of $Y^{(t)}$. To this end, for $x \in \mathbb{R}^n$ define $\text{sign}(x) \in \{-1, 1\}^n$ to be the vector with i -th coordinate $\text{sign}(x)_i = 1 \iff x_i > 0$.

The Generator G :

1. Let X_1, \dots, X_t be independent copies of X for a suitable value $t = O(q^{-1} \cdot \log(n/\varepsilon))$
2. Let $Y^{(0)} = 0^n$, and for $j > 0$ define $Y^{(j)} = Y^{(j-1)} + \delta_{Y^{(j-1)}} \odot X^{(j)}$
3. Output $\text{sign}(Y^{(t)})$

Analysis of the random walk

To prove the correctness of the generator G , we will prove that the random walk has three properties:

- (a) Each step introduces little error: For every $f \in \mathcal{F}$ and $j \in [t]$, $|\mathbb{E}[f(Y^{(j)})] - \mathbb{E}[f(Y^{(j+1)})]| \leq \varepsilon$.
- (b) The walk *polarizes* with high probability: $\Pr[\|\delta_{Y^{(t)}}\|_\infty \leq \varepsilon/n] \geq 1 - \varepsilon$.
- (c) The final rounding operation introduces little error: For every $f \in \mathcal{F}$, conditioned on polarization, $|f(Y^{(t)}) - f(\text{sign}(Y^{(t)}))| \leq \varepsilon$.

We prove these properties in the next three lemmas.

Lemma 5.1.10 (Steps incur small error). *Let \mathcal{F} be a family of functions $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ that is closed under restrictions, and suppose $X \in [-1, 1]^n$ fools \mathcal{F} with error ε . Then for every $f \in \mathcal{F}$ and $y \in [-1, 1]^n$,*

$$|f(y) - \mathbb{E}[f(y + \delta_y \odot X)]| \leq \varepsilon.$$

In particular, for every $j \in [t]$,

$$\left| \mathbb{E}[f(Y^{(j)})] - \mathbb{E}[f(Y^{(j+1)})] \right| \leq \varepsilon.$$

Proof. Let $y \in [-1, 1]^n$ be fixed. Sample a restriction $R \in \{-1, 1, \star\}^n$, independent of X , where the coordinates of R are independent and distributed as follows:

$$R_i = \begin{cases} \text{sign}(y_i) & \text{with probability } |y_i| \\ \star & \text{with probability } 1 - |y_i|. \end{cases}$$

We can extend the composition operation $R \circ x$ to the case of a vector $x \in [-1, 1]^n$ in the natural way: we use x to fill in the \star coordinates of R (see Equation (5.1)). That way, for each coordinate $i \in [n]$, we have

$$\mathbb{E}_R[(R \circ x)_i] = |y_i| \cdot \text{sign}(y_i) + (1 - |y_i|) \cdot x_i = y_i + (1 - |y_i|) \cdot x_i,$$

and hence overall,

$$\mathbb{E}_R[R \circ x] = y + \delta_y \odot x.$$

By Fact 5.1.3, it follows that

$$\mathbb{E}_R[f|_R(x)] = \mathbb{E}_R[f(R \circ x)] = f\left(\mathbb{E}_R[R \circ x]\right) = f(y + \delta_y \odot x).$$

Consequently,

$$\left|f(y) - \mathbb{E}_X[f(y + \delta_y \odot X)]\right| = \left|\mathbb{E}_R[f|_R(0^n)] - \mathbb{E}_{R,X}[f|_R(X)]\right| \leq \mathbb{E}_R\left[\left|f|_R(0^n) - \mathbb{E}_X[f|_R(X)]\right|\right] \leq \varepsilon,$$

where the last step uses the fact that \mathcal{F} is closed under restriction, hence X fools $f|_R$ with error ε for every fixing of R . \square

Next, we will show that the random walk above converges to $\{-1, 1\}^n$ quickly. For this argument, we assume that X is both q -noticeable for a large enough $q > 0$ and *symmetric* as defined below.

Definition 5.1.11 (Symmetric random variables). *Let X be a random variable distributed over $[-1, 1]^n$. We say that X is symmetric if for every $x \in [-1, 1]^n$, we have $\Pr[X = x] = \Pr[X = -x]$.*

We can justify the symmetry assumption as follows. Starting from an arbitrary q -noticeable fractional ε -PRG G_{frac} for \mathcal{F} with seed length s , we can define a new q -noticeable fractional PRG with seed length $s + 1$ by the formula

$$G'_{\text{frac}}(x, b) = (-1)^b \cdot G_{\text{frac}}(x).$$

The distribution $G'_{\text{frac}}(U_{s+1})$ is symmetric, and because \mathcal{F} is closed under shifts, $G'_{\text{frac}}(U_{s+1})$ still fools \mathcal{F} with error ε . (This is the only place where we use the assumption that \mathcal{F} is closed under shifts.)

The symmetry assumption is helpful because of the following lemma concerning the case $n = 1$.

Lemma 5.1.12. *Let $X \in [-1, 1]$ be a symmetric q -noticeable random variable. Then*

$$\mathbb{E}\left[\sqrt{1 - X}\right] \leq 1 - q/8.$$

In their original paper, Chattopadhyay, Hatami, Hosseini, and Lovett observed that Lemma 5.1.12 follows immediately from the Taylor expansion of the function $\sqrt{1 - x}$ [CHHL19]. We present an alternative argument below.

Proof of Lemma 5.1.12. Let $Y = |X|$, and sample $Z \in \{\pm 1\}$ independently of X . Then the product YZ is distributed identically to X . Furthermore, for each fixed value $y \in [0, 1]$, we have

$$\left(\mathbb{E}\left[\sqrt{1 - yZ}\right]\right)^2 = \left(\frac{\sqrt{1 - y} + \sqrt{1 + y}}{2}\right)^2 = \frac{1 + \sqrt{1 - y^2}}{2} \leq 1 - \frac{y^2}{4}.$$

Therefore,

$$\mathbb{E}\left[\sqrt{1 - X}\right] = \mathbb{E}_Y\left[\mathbb{E}_Z\left[\sqrt{1 - YZ}\right]\right] \leq \mathbb{E}_Y\left[\sqrt{1 - Y^2/4}\right] \leq \mathbb{E}_Y[1 - Y^2/8] \leq 1 - q/8. \quad \square$$

Next, let us use [Lemma 5.1.12](#) to show that coordinate-wise polarization happens with high probability. Indeed, looking ahead, the probability will be high enough to allow a union bound over all coordinates.

Lemma 5.1.13 (Polarization). *Let $A^{(1)}, \dots, A^{(t)} \in [-1, 1]$ be independent symmetric q -noticeable random variables. Define $B^{(0)} = 0$, and for $j > 0$ define*

$$B^{(j)} = B^{(j-1)} + (1 - |B^{(j-1)}|) \cdot A^{(j)}. \quad (5.2)$$

Then $\Pr[1 - |B^{(t)}| \geq e^{-tq/8}] \leq e^{-tq/16}$.

Proof. What happens to the distance $1 - |B^{(j)}|$ when we apply the update rule given in [Equation \(5.2\)](#)? If $\text{sign}(A^{(j)}) = \text{sign}(B^{(j-1)})$ (the “good case”), the distance decreases by a factor of $1 - |A^{(j)}|$. If $\text{sign}(A^{(j)}) \neq \text{sign}(B^{(j-1)})$ (the “bad case”), the distance might increase, but at most it increases by a factor of $1 + |A^{(j)}|$. Either way, for $j > 0$, we have

$$1 - |B^{(j)}| \leq (1 - |B^{(j-1)}|) \cdot (1 - A^{(j)} \cdot \text{sign}(B^{(j-1)})).$$

We have assumed that $A^{(1)}, \dots, A^{(j-1)}$ are symmetric. It follows that $B^{(j-1)}$ is also symmetric. Therefore, $|B^{(j-1)}|$ and $A^{(j)} \cdot \text{sign}(B^{(j-1)})$ are independent. As a consequence,

$$\mathbb{E} \left[\sqrt{1 - |B^{(j)}|} \right] \leq \mathbb{E} \left[\sqrt{1 - |B^{(j-1)}|} \right] \cdot \mathbb{E} \left[\sqrt{1 - A^{(j)} \cdot \text{sign}(B^{(j-1)})} \right].$$

The random variable $A^{(j)} \cdot \text{sign}(B^{(j-1)})$ is symmetric and q -noticeable, so we may apply [Lemma 5.1.12](#), giving us

$$\mathbb{E} \left[\sqrt{1 - |B^{(j)}|} \right] \leq \mathbb{E} \left[\sqrt{1 - |B^{(j-1)}|} \right] \cdot (1 - q/8).$$

By induction, this implies that

$$\mathbb{E} \left[\sqrt{1 - |B^{(t)}|} \right] \leq (1 - q/8)^t \leq e^{-qt/8}.$$

The lemma follows by Markov’s inequality. □

Now we show that the final rounding step does not introduce too much error.

Lemma 5.1.14 (Rounding Error). *Let $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ be a function, and extend it to the domain $[-1, 1]^n$ via the Fourier expansion. For every $y \in [-1, 1]^n$,*

$$|f(y) - f(\text{sign}(y))| \leq \sum_{i=1}^n (1 - |y_i|) \leq n \cdot \|\delta_y\|_\infty.$$

Proof. We have

$$\begin{aligned} |f(y) - f(\text{sign}(y))| &= |\mathbb{E}[f(\Pi_y)] - f(\text{sign}(y))| && \text{(Fact 5.1.3)} \\ &\leq 2 \cdot \Pr[\Pi_y \neq \text{sign}(y)] && \text{(since } \|f\|_\infty \leq 1) \\ &\leq 2 \cdot \sum_{i=1}^n \frac{1 - |y_i|}{2}, \end{aligned}$$

where the final inequality follows by the union bound and the observation that the marginal distributions of Π_y are given by

$$(\Pi_y)_i = \begin{cases} \text{sign}(y_i) & \text{with probability } \frac{1+|y_i|}{2} \\ -\text{sign}(y_i) & \text{with probability } \frac{1-|y_i|}{2}. \end{cases} \quad \square$$

We can now analyze G and complete the proof of [Theorem 5.1.8](#). The output of the generator G is $\text{sign}(Y^{(t)})$ for $t = 16 \log(n/\varepsilon)/q$. The seed for G is determined by t independent samples from the fractional generator, and hence has seed-length $ts = O(s \log(n/\varepsilon)/q)$. Let \mathbf{E} denote the event that $\|\delta_{Y^{(t)}}\|_\infty \leq e^{-tq/8} \leq \varepsilon/n$. Then, we can bound the error of the generator $\text{sign}(Y^{(t)})$ as follows:

$$\begin{aligned}
& |\mathbb{E}[f] - \mathbb{E}[f(\text{sign}(Y^{(t)}))]| \\
& \leq |\mathbb{E}[f(\text{sign}(Y^{(t)}))] - \mathbb{E}[f(Y^{(t)})]| + \sum_{j=1}^t |\mathbb{E}[f(Y^{(j)})] - \mathbb{E}[f(Y^{(j-1)})]| \\
& \leq |\mathbb{E}[f(\text{sign}(Y^{(t)})) - f(Y^{(t)})]| + \varepsilon t && \text{(by [Lemma 5.1.10](#))} \\
& \leq |\mathbb{E}[f(\text{sign}(Y^{(t)})) - f(Y^{(t)}) \mid \mathbf{E}]| + 2 \Pr[\mathbf{E}] + \varepsilon t \\
& \leq |\mathbb{E}[f(\text{sign}(Y^{(t)})) - f(Y^{(t)}) \mid \mathbf{E}]| + 2n \cdot e^{-tq/16} + \varepsilon t && \text{(by [Lemma 5.1.13](#))} \\
& \leq |\mathbb{E}[f(\text{sign}(Y^{(t)})) - f(Y^{(t)}) \mid \mathbf{E}]| + \varepsilon(t+2) \\
& \leq (t+2)\varepsilon && \text{(by [Lemma 5.1.14](#))} \\
& \leq O(\varepsilon \log(n/\varepsilon)/q).
\end{aligned}$$

5.1.4 A better reduction for the low-error regime

In this section, we present a more refined reduction, showing how to convert a fractional PRG into a standard PRG with slightly better parameters than what was achieved by Chattopadhyay, Hatami, Hosseini, and Lovett [[CHHL19](#)] ([Theorem 5.1.8](#)).

Theorem 5.1.15 (Fractional PRG \implies Standard PRG, refined version). *Suppose that \mathcal{F} is a family of functions $f: \{-1, 1\}^n \rightarrow \{-1, 1\}$ that is closed under restrictions and shifts. Assume that there exists a q -noticeable fractional PRG for \mathcal{F} with error ε and seed length s . Then there exists an explicit PRG for \mathcal{F} with seed length $O((s + \log(1/\varepsilon)) \cdot q^{-1} \cdot \log n)$ and error $O(\varepsilon \cdot q^{-1} \cdot \log n)$.*

For comparison, recall that in the conclusion of [Theorem 5.1.8](#), the seed length is $O(s \cdot q^{-1} \cdot \log(n/\varepsilon))$ and the error is $O(\varepsilon \cdot q^{-1} \cdot \log(n/\varepsilon))$. We typically have $s \geq \log(1/\varepsilon)$, so the parameters of [Theorem 5.1.15](#) are superior in the regime $\varepsilon < n^{-\omega(1)}$.

The idea behind the improvement is to replace the trivial rounding step. Instead of taking $O(q^{-1} \cdot \log(n/\varepsilon))$ steps of the random walk and arguing that *all* the coordinates of $Y^{(t)}$ are well-polarized (i.e., close to $\{-1, 1\}$), we will take only $O(q^{-1} \cdot \log n)$ steps of the random walk and argue that *most* of the coordinates of $Y^{(t)}$ are well-polarized. Then, we will show how to approximately sample from the mostly-polarized product distribution $\Pi_{Y^{(t)}}$.

More precisely, our notion of being “mostly polarized” is that when we sample from $\Pi_{Y^{(t)}}$, with high probability, we get a vector that only disagrees with $\text{sign}(Y^{(t)})$ in a few coordinates:

Definition 5.1.16 (Polarization). *Let $y \in [-1, 1]^n$, let $k \in \mathbb{N}$, and let $\delta > 0$. We say that y is (k, δ) -polarized if*

$$\Pr[\Delta(\Pi_y, \text{sign}(y)) \leq k] \geq 1 - \delta,$$

where Δ denotes Hamming distance.

We now show that $O(q^{-1} \cdot \log n)$ steps of the random walk suffice to achieve $(2k, \varepsilon)$ -polarization where $k = O(\log(1/\varepsilon)/\log n)$. For this argument, we assume that the coordinates of the output distribution X of the fractional PRG are k -wise independent. To justify this assumption, observe that we can replace X with $X \odot X'$, where $X' \in \{\pm 1\}^n$ is a k -wise independent distribution. Since \mathcal{F} is closed under shifts, this distribution still fools \mathcal{F} with error ε . This modification only increases

the seed length of the fractional PRG by an additive $O(k \log n) = O(\log(1/\varepsilon))$ bits; this is the reason for the $s + \log(1/\varepsilon)$ term in the conclusion of [Theorem 5.1.15](#).

Lemma 5.1.17 (Polarization, refined version). *Let $k = \lceil \log(1/\varepsilon)/\log n \rceil$ and assume that the coordinates of X are k -wise independent.¹ There exists a value $t = O(p^{-1} \log n)$ such that with probability $1 - \varepsilon$, the vector $Y^{(t)}$ is $(2k, \varepsilon)$ -polarized.*

Proof. Let J be the set of coordinates where $Y^{(t)}$ is “poorly polarized,” namely

$$J = \left\{ i \in [n] : 1 - |Y_i^{(t)}| \geq \frac{1}{n^2} \right\}.$$

By [Lemma 5.1.13](#), there is a choice of $t = O(p^{-1} \log n)$ such that for each coordinate $i \in [n]$, we have $\Pr[i \in J] \leq 1/n^2$. Therefore, for any set $S \subseteq [n]$ with $|S| = k$, we have

$$\Pr[S \subseteq J] \leq n^{-2k}.$$

Thus by a union bound,

$$\Pr[|J| \geq k] \leq \binom{n}{k} \cdot n^{-2k} \leq n^{-k} \leq \varepsilon.$$

Now, fix any $y \in [-1, 1]^n$ such that $|\{i \in [n] : 1 - |y_i| \geq 1/n^2\}| < k$. Every such y is $(2k, \varepsilon)$ -polarized, because

$$\Pr[\Delta(\Pi_y, \text{sign}(y)) \geq 2k] \leq \binom{n-k}{k} \cdot \left(\frac{1}{2n^2}\right)^k \leq n^{-k} \leq \varepsilon. \quad \square$$

Next, as outlined before, we show that when y is $(2k, \varepsilon)$ -polarized, we can approximately sample from Π_y . We use a naïve “brute force” approach.

Lemma 5.1.18 (Approximately sampling from well-polarized product distributions). *Let $y \in [-1, 1]^n$, let $k \in \mathbb{N}$, and let $\delta > 0$. There is a randomized algorithm `Sample` such that `Sample(y, k, δ)` outputs a string in $\{-1, 1\}^n$, and if y is (k, δ) -polarized, then the output distribution `Sample(y, k, δ)` is within total variation distance $O(\delta)$ of Π_y . Furthermore, `Sample(y, k, δ)` runs in time $\text{poly}(n^k, 1/\delta)$ and it uses $O(k \log n + \log(1/\delta))$ truly random bits.*

Proof. Let $z = \frac{1}{2}\delta_y \in [0, 1/2]^n$. Let D be the product distribution over $\{0, 1\}^n$ such that $\mathbb{E}[D] = z$. Note that if we could sample $x \sim D$, then we could sample Π_y by outputting the vector with i -th coordinate $(-1)^{x_i} \cdot \text{sign}(y_i)$. Thus, it suffices to show how to efficiently sample a vector from $\{0, 1\}^n$ with a distribution close to D in total variation distance. To achieve this, it is helpful to think of a perfect sample from D as being produced by the following process.

Perfectly sampling D :

1. Pick $\rho \in [0, 1]$ uniformly at random.
2. Initialize $\mu := 0$
3. For $e \in \{0, 1\}^n$:
 - 3.1 $\mu := \mu + \prod_{i:e_i=0} z_i \prod_{i:e_i=1} (1 - z_i)$
 - 3.2 If $\mu \geq \rho$ then halt and output e
4. Output 0^n

¹In fact, it suffices for the coordinates to be ε^2 -almost k -wise independent.

To *efficiently* sample from D (approximately), we make two changes. First, in the “for loop,” we only iterate over e that satisfy $\sum_{i=1}^n e_i \leq k$. By the assumption of (k, δ) -polarization, this change only introduces total variation error at most δ . Second, we discretize ρ . That is, we sample $\rho \in \{0, \alpha, 2\alpha, \dots, 1\}$ uniformly at random, where $\alpha = \delta \cdot n^{-k}$. The additional total variation error introduced by this second change is at most $\binom{n}{k} \cdot \alpha \leq \delta$. \square

Given [Lemmas 5.1.17](#) and [5.1.18](#), it follows that the PRG below proves [Theorem 5.1.15](#), assuming that X is symmetric and its coordinates are k -wise independent.

The Generator G' :

1. Let $t = O(q^{-1} \cdot \log(n))$ and $k = \lceil \log(1/\varepsilon) / \log n \rceil$
2. Let $X^{(1)}, \dots, X^{(t)}$ be independent copies of X
3. Let $Y^{(1)} = 0^n$, and for $j > 0$ define $Y^{(j)} = Y^{(j-1)} + \delta_{Y^{(j-1)}} \odot X^{(j)}$
4. Output $\text{Sample}(Y^{(t)}, 2k, \varepsilon)$

5.2 Analysis technique: Fourier growth bounds

Let \mathcal{F} be a class of functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$ that we wish to fool, and let $\mathcal{F}_{\text{simp}}$ be a class of “simpler” functions that we know how to fool. In the previous section, we considered the case that every $f \in \mathcal{F}$ simplifies to $\mathcal{F}_{\text{simp}}$ with high probability under random restrictions, i.e., for some $p, \delta > 0$, we have

$$\Pr[f|_{\mathcal{R}_p} \in \mathcal{F}_{\text{simp}}] \geq 1 - \delta. \tag{5.3}$$

We presented the “polarizing random walks” framework, which shows that under this assumption, we can construct a PRG for \mathcal{F} .

In this section, we consider a more general setting, which is when \mathcal{F} “simplifies on average” under restrictions. We explain the meaning of this condition in [Section 5.2.1](#). Then, in [Section 5.2.2](#), we present an example of this condition – we show that bounded-width regular ROBPs satisfy such an “average-case” simplification-under-restrictions lemma. Finally, in [Section 5.2.3](#), we show that the polarizing random walks framework still works under this weaker assumption, and consequently we get a PRG for the model of bounded-width “arbitrary-order permutation ROBPs.”

5.2.1 The noise operator and simplification on average

The notion of “simplification on average” is based on the *noise operator*.

Definition 5.2.1 (Noise operator). *Let $f: \{0, 1\}^n \rightarrow \mathbb{R}$ and $p > 0$. We define $T_p f: \{0, 1\}^n \rightarrow \mathbb{R}$ by the equation*

$$T_p f(x) = \mathbb{E}[f|_{\mathcal{R}_p}(x)].$$

T_p is called the “noise operator” with parameter p , because for each bit of x , with probability $1 - p$, we replace the bit with a fresh random bit (“noise”). Intuitively, $T_p f$ is a “smoothed out” version of f , and smaller values of p correspond to more smoothing out.

We say that f *simplifies on average* under the random restriction \mathcal{R}_p if the function $T_p f$ lies in some “simpler” class $\mathcal{F}_{\text{simp}}$. For example, let \mathcal{F} be the class of parity functions. When we apply a

random restriction \mathcal{R}_p , with high probability, no meaningful simplification occurs: the restriction of a parity function is another parity function (or its complement). However, parity functions *do* drastically simplify *on average* over restrictions. Indeed, if f is a parity function on k bits, then $T_p f$ is approximated by the constant $1/2$ function to within pointwise error p^{-k} .

For a more interesting example, let us return to the model of *bounded-width regular ROBPs*, which we studied previously in [Section 3.3](#). These programs can compute parity functions, so once again, they do not meaningfully simplify under a typical *individual* restriction. However, we will show that these programs simplify on average under restrictions. Specifically, we will show that $T_p f$ is fooled by almost k -wise uniform distributions, where p and k are suitable parameters and f is any bounded-width regular ROBP. Our approach for proving this simplification-under-restrictions lemma is to bound the *Fourier growth* of f , discussed next.

5.2.2 Fourier growth bounds for regular ROBPs

Recall the Fourier L_1 bound from [Section 2.3](#), which is a simple Fourier-analytic way of measuring the “complexity” of a Boolean function. Fourier growth is a more refined complexity measure which takes into account the *degree* of Fourier coefficients. Specifically:

Definition 5.2.2 (Functions with bounded Fourier growth). *For $a, b \geq 1$, denote by $\mathcal{L}_1^n(a, b)$ the family of all n -variate Boolean functions $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ that satisfy*

$$\sum_{\substack{S \subseteq [n] \\ |S|=d}} |\hat{f}(S)| \leq a \cdot b^d,$$

for every $d \in [n]$. Define $\mathcal{L}_1(a, b) = \bigcup_{n \in \mathbb{N}} \mathcal{L}_1^n(a, b)$.

Remark 5.2.3 (Fourier L_2 tail bounds). *One can similarly define $\mathcal{L}_2^n(a, b)$ to be the family of all n -variate Boolean functions $f : \{-1, 1\}^n \rightarrow \{-1, 1\}$ that satisfy*

$$\sum_{\substack{S \subseteq [n] \\ |S|=d}} |\hat{f}(S)|^2 \leq a \cdot 2^{-d/b},$$

for every $d \in [n]$. Tal showed that $\mathcal{L}_2(a, b) \subseteq \mathcal{L}_1(a, O(b))$ [[Tal17](#)]. The simple example of the PARITY function (i.e., $\prod_{i \in [n]} x_i$) shows that the reverse is not true. In other words, having bounded L_1 Fourier growth is a weaker assumption than having bounded L_2 Fourier tails.

Reingold, Steinke, and Vadhan were the first to prove a Fourier growth bound for regular ROBPs [[RSV13](#)]. Later, building on their work and work by Chattopadhyay, Hatami, Reingold, and Tal [[CHRT18](#)], Lee, Pyne, and Vadhan improved the bound [[LPV22](#)]. In this section, we will present the proof of the latter bound.

Theorem 5.2.4 (Fourier growth of regular ROBPs [[LPV22](#)]). *If f is a width- w standard-order² regular ROBP, then $f \in \mathcal{L}_1(1, w - 1)$. That is, for every $d \geq 0$,*

$$\sum_{|S|=d} |\hat{f}(S)| \leq (w - 1)^d.$$

After we are done proving [Theorem 5.2.4](#), we will show that Fourier growth bounds imply simplification on average under restrictions.

²The theorem holds more generally for the “arbitrary-order” model, in which the ROBP reads the variables in an arbitrary permuted order (note that we still assume the ROBP is oblivious). The reason is that the quantity $\sum_{|S|=d} |\hat{f}(S)|$ is invariant under variable permutations.

Level 1 Fourier coefficients

The first step of the proof of [Theorem 5.2.4](#) is to bound the level-1 Fourier coefficients of f . As a shorthand, we write $\widehat{f}(i)$ rather than $\widehat{f}(\{i\})$. We will prove the following.

Lemma 5.2.5 (Bound on level-1 Fourier coefficients of regular ROBPs). *Let f be a width- w length- n standard-order regular ROBP. Then*

$$\sum_{i=1}^n |\widehat{f}(i)| \leq \mathbb{E}[f] \cdot (w - 1).$$

Proof. Let m be the number of rejecting vertices in the final layer, i.e., $m = w - |V_{\text{accept}}|$. We will show by induction on n that

$$\sum_{i=1}^n |\widehat{f}(i)| \leq \mathbb{E}[f] \cdot m. \quad (5.4)$$

The lemma will follow, because $m \leq w$ (and if $m = w$, then $f \equiv 0$ and the lemma is trivial).

In the base case $n = 0$, [Equation \(5.4\)](#) is trivial, so assume that $n > 0$. Let V_0, \dots, V_n be the layers of f . Partition $V_{n-1} = R \cup S \cup T$ based on the number of accepting edges, i.e.,

$$\begin{aligned} R &= \{v \in V_{n-1} : \mathbb{E}[f_{v \rightarrow}] = 0\} \\ S &= \{v \in V_{n-1} : \mathbb{E}[f_{v \rightarrow}] = 1/2\} \\ T &= \{v \in V_{n-1} : \mathbb{E}[f_{v \rightarrow}] = 1\}. \end{aligned}$$

Observe that $m = |R| + \frac{1}{2}|S|$ because f is regular. For each $i < n$, we have

$$\begin{aligned} \widehat{f}(i) &= \mathbb{E}_{x \sim U_n} [f(x) \cdot (-1)^{x_i}] = \mathbb{E}_{x_1, \dots, x_{n-1}} \left[(-1)^{x_i} \cdot \mathbb{E}_{x_n} [f(x)] \right] \\ &= \mathbb{E}_{x_1, \dots, x_{n-1}} \left[(-1)^{x_i} \cdot \left(f_{\rightarrow T}(x) + \frac{1}{2} f_{\rightarrow S}(x) \right) \right] \\ &= \widehat{f_{\rightarrow T}}(i) + \frac{1}{2} \widehat{f_{\rightarrow S}}(i). \end{aligned}$$

Therefore, if we write $p_{\rightarrow A}$ as a shorthand for the probability of visiting a vertex in $A \subseteq V_{n-1}$ (namely, $p_{\rightarrow A} = \mathbb{E}[f_{\rightarrow A}]$), then we have

$$\begin{aligned} \sum_{i=1}^{n-1} |\widehat{f}(i)| &\leq \frac{1}{2} \sum_{i=1}^{n-1} |\widehat{f_{\rightarrow T}}(i)| + \frac{1}{2} \sum_{i=1}^{n-1} |\widehat{f_{\rightarrow T}}(i) + \widehat{f_{\rightarrow S}}(i)| \\ &= \frac{1}{2} \sum_{i=1}^{n-1} |\widehat{f_{\rightarrow T}}(i)| + \frac{1}{2} \sum_{i=1}^{n-1} |\widehat{f_{\rightarrow S \cup T}}(i)| \\ &\leq \frac{1}{2} |R \cup S| \cdot p_{\rightarrow T} + \frac{1}{2} |R| \cdot (p_{\rightarrow S \cup T}) \quad (\text{Induction}) \\ &= m \cdot p_{\rightarrow T} + |R| \cdot \frac{p_{\rightarrow S}}{2}. \end{aligned}$$

Meanwhile, at $i = n$, we have

$$|\widehat{f}(n)| \leq \mathbb{E}_{x_1, \dots, x_{n-1}} \left[\left| \mathbb{E}_{x_n} [(-1)^{x_n} \cdot f(x)] \right| \right] = \frac{p_{\rightarrow S}}{2} \leq \frac{|S|}{2} \cdot \frac{p_{\rightarrow S}}{2},$$

because the regularity of f implies that $|S|$ is even. Therefore, overall,

$$\begin{aligned} \sum_{i=1}^n |\widehat{f}(i)| &\leq m \cdot p_{\rightarrow T} + \left(|R| + \frac{|S|}{2} \right) \cdot \frac{p_{\rightarrow S}}{2} \\ &= m \cdot \left(p_{\rightarrow T} + \frac{p_{\rightarrow S}}{2} \right) \\ &= m \cdot \mathbb{E}[f]. \end{aligned} \quad \square$$

Higher-level Fourier coefficients

To bound the higher-level Fourier coefficients, we rely on a notion of *local monotonicity* [BV10; CHRT18]. For every vertex v of an ROBP f , define $p_{v \rightarrow} = \mathbb{E}[f_{v \rightarrow}]$.

Definition 5.2.6 (Local Monotonicity). *Let f be a standard-order ROBP with layers V_0, \dots, V_n . We say that f is locally monotone if for each $i \in [n]$ and each vertex $u \in V_{i-1}$, if we let (u, v) be the outgoing 0-edge and (u, s) be the outgoing 1-edge, then $p_{s \rightarrow} \geq p_{v \rightarrow}$.*

It is easy to see that if f is a locally monotone ROBP, then for every i , we have $\widehat{f}(i) \leq 0$.³ We observe that every ROBP can be “locally monotone” by relabeling its edges, and so local monotonicity is a property of the labeling and not the structure of the graph.

Observation 5.2.7 (Local Monotonization [CHRT18]). *Let f be a standard-order ROBP with layers V_0, \dots, V_n . There exists a locally monotone standard-order ROBP f' obtained by relabeling edges of f . Furthermore, for every $i \in [n]$ and every $v \in V_{i-1}$,*

$$\widehat{f'_{v \rightarrow}}(i) = - \left| \widehat{f_{v \rightarrow}}(i) \right|.$$

Proof. First order the vertices in each layer according to $p_{v \rightarrow}$, breaking ties according to a predetermined fixed ordering. We start from the layer V_i for $i = n$ and move backwards. For every vertex $v \in V_i$ we relabel its outgoing edges if they do not satisfy the local monotonicity condition. Note that for each vertex v , the acceptance probability $p_{v \rightarrow}$ remains unchanged under this relabeling, and hence the ordering of the vertices within each layer remains the same. Furthermore, swapping the labels of the outgoing edges of a vertex $v \in V_{i-1}$ flips the sign of the Fourier coefficient $\widehat{f_{v \rightarrow}}(i)$ so that it is nonpositive. \square

Note that if f is regular, then so is the local monotonization f' . Given [Observation 5.2.7](#), we are prepared to bound the higher-level Fourier coefficients of a regular ROBP, using an inductive argument due to Chattopadhyay, Hatami, Reingold, and Tal [CHRT18].

Proof of [Theorem 5.2.4](#). We will show by induction on d that

$$\sum_{|S|=d} |\widehat{f}(S)| \leq (w-1)^d \cdot \mathbb{E}[f].$$

³In the literature, the reverse inequality is claimed [CHRT18; LPV22], but this seems to be a mistake.

Lemma 5.2.5 is the base case $d = 1$. For the inductive step, let the layers of f be V_0, \dots, V_n . Then

$$\begin{aligned}
\sum_{|S|=d+1} |\widehat{f}(S)| &= \sum_{i=1}^n \sum_{\substack{T \subseteq [i-1] \\ |T|=d}} |\widehat{f}(T \cup \{i\})| = \sum_{i=1}^n \sum_{\substack{T \subseteq [i-1] \\ |T|=d}} \left| \sum_{v \in V_{i-1}} \widehat{f}_{\rightarrow v}(T) \cdot \widehat{f}_{v \rightarrow}(i) \right| \\
&\leq \sum_{i=1}^n \sum_{\substack{T \subseteq [i-1] \\ |T|=d}} \sum_{v \in V_{i-1}} |\widehat{f}_{\rightarrow v}(T)| \cdot |\widehat{f}_{v \rightarrow}(i)| \\
&= \sum_{i=1}^n \sum_{v \in V_{i-1}} \left(\sum_{\substack{T \subseteq [i-1] \\ |T|=d}} |\widehat{f}_{\rightarrow v}(T)| \right) \cdot |\widehat{f}_{v \rightarrow}(i)| \\
&\leq (w-1)^d \cdot \sum_{i=1}^n \sum_{v \in V_{i-1}} \mathbb{E}[f_{\rightarrow v}] \cdot |\widehat{f}_{v \rightarrow}(i)|
\end{aligned}$$

by induction. Now, to get rid of the absolute value signs, let f' be the local monotonization of f from Observation 5.2.7. Then continuing, we have

$$\begin{aligned}
\sum_{|S|=d+1} |\widehat{f}(S)| &\leq (w-1)^d \cdot \sum_{i=1}^n \sum_{v \in V_{i-1}} \mathbb{E}[f_{\rightarrow v}] \cdot |\widehat{f}_{v \rightarrow}(i)| \\
&= (w-1)^d \cdot \sum_{i=1}^n \sum_{v \in V_{i-1}} \mathbb{E}[f'_{\rightarrow v}] \cdot (-\widehat{f}_{v \rightarrow}(i)) \\
&= (w-1)^d \cdot \sum_{i=1}^n - \mathbb{E}_{x \sim U_n} \left[\sum_{v \in V_{i-1}} f'_{\rightarrow v}(x) \cdot f'_{v \rightarrow}(x) \cdot (-1)^{x_i} \right] \\
&= (w-1)^d \cdot \sum_{i=1}^n -\widehat{f}'(i) \\
&\leq (w-1)^d \cdot (w-1) \cdot \mathbb{E}[f'] && \text{(Lemma 5.2.5)} \\
&= (w-1)^{d+1} \cdot \mathbb{E}[f]. && \square
\end{aligned}$$

Fourier growth bounds imply simplification on average

So far, we have shown that regular ROBPs have bounded Fourier growth. Next, we show that in general, functions with bounded Fourier growth simplify on average under restrictions, to the point that they can be fooled by k -wise small-bias distributions.

Proposition 5.2.8 (Bounded Fourier growth \implies simplification on average under restrictions). *Let $a, b \geq 1$ and $f \in \mathcal{L}_1^n(a, b)$. Let $\varepsilon \in (0, 1)$, and let $X \in \{-1, 1\}^n$ be k -wise δ -biased,⁴ where*

$$\delta = \varepsilon / (2a) \quad \text{and} \quad k = \lceil \log(2a/\varepsilon) \rceil.$$

Let $p = 1/(2b)$. Then X fools $T_p f$ with error ε .

⁴Since we are working over ± 1 , the meaning of “ k -wise δ -biased” is that for every nonempty set $S \subseteq [n]$ with $|S| \leq k$, we have $|\mathbb{E}[\prod_{i \in S} X_i]| \leq \delta$.

Proof. The Fourier expansion of $T_p f$ is given by

$$\begin{aligned} T_p f(x) &= \mathbb{E}_{R \sim R_p} [f(R \circ x)] = \sum_{S \subseteq [n]} \widehat{f}(S) \mathbb{E}_{R \sim R_p} [\chi_S(R \circ x)] \\ &= \sum_{S \subseteq [n]} \widehat{f}(S) \cdot p^{|S|} \cdot \chi_S(x). \end{aligned}$$

(Informally, the noise operator T_p “attenuates” the Fourier coefficients of f .) Noting that $\widehat{T_p f}(\emptyset) = \mathbb{E}[T_p f]$, we have

$$|\mathbb{E}[T_p f(X)] - \mathbb{E}[T_p f]| = \left| \sum_{\substack{S \subseteq [n] \\ S \neq \emptyset}} p^{|S|} \widehat{f}(S) \mathbb{E} \left[\prod_{i \in S} X_i \right] \right| \leq \sum_{\substack{S \subseteq [n] \\ S \neq \emptyset}} p^{|S|} |\widehat{f}(S)| \cdot \left| \mathbb{E} \left[\prod_{i \in S} X_i \right] \right|.$$

When $|S| \leq k$, we have $|\mathbb{E}[\prod_{i \in S} X_i]| \leq \delta$. When $|S| > k$, we use the trivial bound $|\mathbb{E}[\prod_{i \in S} X_i]| \leq 1$. Plugging these bounds into the above inequality, we get

$$|\mathbb{E}[T_p f(X)] - \mathbb{E}[T_p f]| \leq \delta \cdot a \cdot \sum_{d=1}^k (pb)^d + a \cdot \sum_{d=k+1}^n (pb)^d \leq \delta a + 2^{-k} a \leq \varepsilon. \quad \square$$

Thus, in particular, bounded-width standard-order regular ROBPs simplify on average under restrictions.

5.2.3 Using Fourier growth bounds to obtain PRGs

Using the analysis in the previous section, let us now obtain a PRG for ROBPs. We already presented some PRGs for *standard-order* ROBPs in [Chapter 3](#), such as the INW PRG. However, those PRGs rely heavily on the fact that the order of the input variables of the ROBP is known to the generator. A more challenging scenario is when we wish to fool functions of the form

$$f(x) = g(x_{\pi(1)}, \dots, x_{\pi(n)}),$$

where g is a width- w length- n standard-order ROBP and π is an *unknown* permutation of the coordinates $[n]$. We refer to such a function f as a width- w length- n *arbitrary-order ROBP* (see [Figure 5.2](#)). One motivation for fooling arbitrary-order ROBPs is that they can simulate other interesting classes of tests, such as read-once formulas [\[BPW11\]](#). In this section, we will focus on a subclass of ROBPs called *permutation ROBPs*.

Definition 5.2.9 (Permutation ROBPs). *Let f be a length- n arbitrary-order ROBP with layers V_0, \dots, V_n . We say that f is a permutation ROBP if for every $i \in [n]$ and $v \in V_i$, there are exactly two incoming edges to v (regularity), and those two edges have distinct labels (one is labeled 0 and the other is labeled 1).*

In [Section 3.3](#), we saw the BRRY generator, which fools constant-width *standard-order* regular ROBPs with seed length $\tilde{O}(\log n)$. We will now show how to design another PRG for constant-width permutation ROBPs, which once again has seed length $\tilde{O}(\log n)$, but this time the PRG works even in the more challenging arbitrary-order setting:




Figure 5.2: Define $f: \{0, 1\}^{2n} \rightarrow \{0, 1\}$ by the formula $f(x) = \bigoplus_{1 \leq i < j \leq n} x_i \cdot x_{n+j}$. This function can be computed by a width-4 arbitrary-order permutation ROBP (the case $n = 2$ is shown above). In contrast, every *standard-order* ROBP computing f has width $2^{\Omega(n)}$. This can be shown by a communication complexity argument, reducing from the inner product mod 2 problem.

Theorem 5.2.10 (PRG for arbitrary-order permutation ROBPs [RSV13; CHHL19; LPV22]). *For every $w, n \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for width- w length- n arbitrary-order permutation ROBPs with seed length $\tilde{O}(w^2 \cdot \log n \cdot \log(1/\varepsilon))$.⁵*

As a reminder, we showed in Section 5.2.2 that permutation ROBPs (and more generally regular ROBPs) simplify on average under restrictions. To prove Theorem 5.2.10, we now observe that such an average-case simplification-under-restrictions lemma is sufficient for the polarizing random walks framework.

Lemma 5.2.11 (Simplification on average implies fractional PRGs). *Let \mathcal{F} be a class of functions $f: \{\pm 1\}^n \rightarrow \mathbb{R}$. Let $p, \varepsilon > 0$, let X be a distribution over $\{-1, 1\}^n$, and assume that for every $f \in \mathcal{F}$, the distribution X fools $T_p f$ with error ε . Then pX is (p^2) -noticeable and fools \mathcal{F} with error ε .*

Proof. The proof is exactly the same as the proof of Lemma 5.1.6, except that we replace the final step with the following:

$$\left| \mathbb{E}_{R, X} [f(R \circ X)] - \mathbb{E}[f] \right| = \left| \mathbb{E}_X [T_p f(X)] - \mathbb{E}[T_p f] \right| \leq \varepsilon. \quad \square$$

Putting everything together gives us our PRG for bounded-width arbitrary-order permutation ROBPs:

Proof of Theorem 5.2.10. Let f be a width- w length- n arbitrary-order permutation ROBP. Such a program f satisfies the Fourier growth bound of Theorem 5.2.4 regardless of the order in which it reads the input variables, because the quantity $\sum_{|S|=d} |\hat{f}(S)|$ is invariant under variable permutations.

⁵The specific seed length in Theorem 5.2.10 does not appear in prior work, but it does follow directly from prior work [RSV13; CHHL19; LPV22; FK18]. In particular, if $\varepsilon > 1/n$, then it follows from Lee, Pyne, and Vadhan's work [LPV22], whereas if $\varepsilon \leq 1/n$, then it follows from Forbes and Kelley's work [FK18]. Our refined polarizing random walks framework (Theorem 5.1.15) gives us a unified proof of Theorem 5.2.10 for all ε .

Let $\gamma = \varepsilon/t$ for a suitable value $t = O(w^2 \log n)$. Let $\delta = \gamma/2$, let $k = \lceil \log(2/\gamma) \rceil$, and let $X \in \{-1, 1\}^n$ be k -wise δ -biased. Let $p = \frac{1}{2 \cdot (w-1)}$. Then by [Theorem 5.2.4](#) and [Proposition 5.2.8](#), X fools $T_p f$ with error γ . Therefore, by [Lemma 5.2.11](#), pX is (p^2) -noticeable and fools f with error γ .

The distribution X can be explicitly sampled using a seed of length $s = O(\log(1/\gamma) + \log \log n)$ (see [Theorem 2.2.7](#)). Furthermore, the class of width- w permutation ROBPs is closed under restrictions and shifts. Therefore, by the refined polarizing random walks framework ([Theorem 5.1.15](#)), there is an explicit PRG for such programs with error $O(\gamma \cdot p^{-2} \cdot \log n) = \varepsilon$ and seed length

$$O(s \cdot p^{-2} \cdot \log n) = O(w^2 \cdot \log n \cdot (\log(w/\varepsilon) + \log \log n)). \quad \square$$

To be clear, the Fourier growth bound ([Theorem 5.2.4](#)) works for all *regular* ROBPs, not merely permutation ROBPs. However, the class of width- w arbitrary-order regular ROBPs is unfortunately not closed under restrictions. It is therefore unclear how to apply the polarizing random walks framework to such programs.

Open Problem 5.2.12 (PRGs for arbitrary-order regular ROBPs). *Design a PRG for constant-width arbitrary-order regular ROBPs with seed length $o(\log^2 n)$.*

Thankfully, the more restricted class consisting of width- w arbitrary-order *permutation* ROBPs is closed under restrictions, a fact which is crucial in the proof of [Theorem 5.2.10](#).

More generally, by the same argument, one can use the polarizing random walks framework to construct a PRG for any class with bounded Fourier growth, provided that the class is closed under restrictions and shifts.⁶

Theorem 5.2.13 (PRG for functions with bounded Fourier growth). *For every n, a, b, ε , there is an explicit PRG G such that if $\mathcal{F} \subseteq \mathcal{L}_1(a, b)$ and \mathcal{F} is closed under restrictions and shifts, then G fools \mathcal{F} with error ε . Furthermore, G has seed length*

$$O(b^2 \cdot \log n \cdot (\log(ab/\varepsilon) + \log \log n)).$$

When Reingold, Steinke, and Vadhan proved their Fourier growth bound for standard-order regular ROBPs, they also conjectured a Fourier growth bound for all standard-order ROBPs (whether regular or not) [[RSV13](#)]. In particular, they conjectured that constant-width standard-order ROBPs are in $\mathcal{L}_1(\text{poly}(n), \text{polylog}(n))$. The width-3 case of this conjecture was proven by Steinke, Vadhan, and Wan [[SVW17](#)], and then the general case was proven by Chattopadhyay, Hatami, Reingold, and Tal [[CHRT18](#)].

Theorem 5.2.14 (Fourier growth bound for ROBPs [[CHRT18](#)]). *Suppose f is a width- w length- n standard-order RBP. Then $f \in \mathcal{L}_1(1, O(\log n)^w)$. That is, for every $d \in [n]$,*

$$\sum_{|S|=d} |\hat{f}(S)| \leq O(\log n)^{wd}.$$

Combining [Theorems 5.2.13](#) and [5.2.14](#) gives a PRG for constant-width *arbitrary-order* ROBPs with seed length $\text{polylog}(n)$. Later ([Section 5.4](#)), we will see a better PRG for this class that will make use of [Theorem 5.2.14](#) in a more sophisticated way.

Motivated by the goal of constructing new PRGs for classes of functions such as $\mathbf{AC}^0[\oplus]$ and \mathbb{F}_2 -polynomials (see [Open Problems 2.4.13](#) and [4.3.1](#)), there has been effort on two fronts with the

⁶One can show that closure under shifts is not necessary. On the other hand, it seems quite challenging to handle classes that are not closed under restrictions.

latter having led to some success: 1. Prove reasonable Fourier tail bounds for the said classes, and 2. Construct fractional PRGs under more relaxed assumptions, for example, Fourier tail bounds only on few levels. Chattopadhyay, Hatami, Lovett, and Tal showed how to construct a fractional PRG using only *second-level* Fourier bounds [CHLT18]. Then, Chattopadhyay, Gaitonde, Lee, Lovett, and Shetty showed that better bounds can be achieved if bounds on higher Fourier levels are available, and interestingly, that fractional PRGs can be achieved even from bounds on $|\sum_{S:|S|=d} \widehat{f}(S)|$ where one can have cancellations, as opposed to L_1 bounds [CGLLS21].

5.3 Fooling \mathbf{AC}^0 via the Ajtai-Wigderson framework

Let \mathcal{F} be a class of functions $f: \{0,1\}^n \rightarrow \{0,1\}$ that we wish to fool. In this section, we revisit the case that functions in \mathcal{F} simplify *with high probability* under restrictions. That is, suppose that for some values $p, \delta > 0$, we have

$$\Pr[f|_{\mathcal{R}_p} \in \mathcal{F}_{\text{simp}}] \geq 1 - \delta, \tag{5.5}$$

and furthermore suppose that we already have an explicit PRG for the “simpler” class $\mathcal{F}_{\text{simp}}$, say with seed length s .

In [Section 5.1](#), we presented the polarizing walks framework, which allows us to construct a PRG for the original class \mathcal{F} with a seed length of roughly $p^{-2} \cdot s \cdot \log n$. In this section, we present an older PRG framework due to Ajtai and Wigderson [AW89]. The Ajtai-Wigderson framework achieves a better seed length, but it requires a stronger initial assumption. Roughly speaking, we will show that if it is possible to “derandomize the choice of \star positions” in [Equation \(5.5\)](#), then we can fool \mathcal{F} with a seed length of approximately $p^{-1} \cdot s \cdot \log n$. The distinction between p^{-2} and p^{-1} is quite important in many cases.

5.3.1 Simplification under partially pseudorandom restrictions

The Ajtai-Wigderson framework is based on *partially pseudorandom restrictions*, meaning that the set of \star positions is pseudorandom, but the assigned values are truly random. To reason about these two components of a restriction separately, we must introduce notation for encoding restrictions using bitstrings. Many different notational approaches have been used for this encoding; unfortunately, it seems inevitable that the notation is somewhat cumbersome and clunky. We will take the approach of defining the following \odot operation.

Definition 5.3.1 (Encoding restrictions). *For $x, y \in \{0,1\}^n$, define $x \odot y \in \{0,1,\star\}^n$ by*

$$(x \odot y)_i = \begin{cases} x_i & \text{if } y_i = 0 \\ \star & \text{if } y_i = 1. \end{cases}$$

(See [Figure 5.3](#).)

The assumption of the Ajtai-Wigderson framework is that we have a simplification-under-restrictions lemma of the form

$$\Pr[f|_{U \odot Z} \in \mathcal{F}_{\text{simp}}] \geq 1 - \delta, \tag{5.6}$$

where the assigned bits $U \in \{0,1\}^n$ is distributed uniformly at random, the star set $Z \in \{0,1\}^n$ is independent of U , and Z can be sampled explicitly with a short seed. For example, for \mathbf{AC}^0 , there is a line of work on proving *derandomized switching lemmas* [AW89; TX13; Tal17; ST19a;




Figure 5.3: In the restriction $x \circledast y$, the string y indicates the \star positions, and the string x assigns values to the non- \star positions.

[Kel21; Lyu22]. By combining Lyu’s “derandomized multi-switching lemma” [Lyu22] with Rossman’s arguments [Ros19], one can prove the following.⁷

Lemma 5.3.2 (Simplification of \mathbf{AC}^0 under a partially pseudorandom restriction [Lyu22; Ros19]). *For every n, m, d, δ , there is a random variable Z over $\{0, 1\}^n$ that can be explicitly sampled using $\tilde{O}(d \cdot \log(mn/\delta))$ truly random bits such that for any depth- d size- m \mathbf{AC}^0 circuit f on n input bits,*

$$\Pr[\text{DT}(f|_{U \circledast Z}) \leq O(\log(md/\delta))] \geq 1 - \delta.$$

Here, U is a uniform random n -bit string independent of Z . Furthermore, for each coordinate i , the “star probability” is given by $\mathbb{E}[Z_i] = \Theta(1/\log m)^{d-1}$.

We will not study the proof of Lemma 5.3.2 here. Instead, we will focus on the logic of constructing a PRG given a statement such as Lemma 5.3.2. The only fact we will use about shallow decision trees is that they can be fooled with a short seed length. Indeed, in its simplest form, the Ajtai-Wigderson framework reduces the problem of constructing PRGs to the problem of proving statements like Equation (5.6) with the following parameters.

Theorem 5.3.3 (Simplification under partially-pseudorandom restrictions \implies PRG [AW89]). *Let \mathcal{F} and $\mathcal{F}_{\text{simp}}$ be classes of functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$. Assume that \mathcal{F} is closed under restrictions. Let Z be a random variable over $\{0, 1\}^n$ that can be explicitly sampled using s truly random bits such that*

$$\forall f \in \mathcal{F}, \Pr[f|_{U \circledast Z} \in \mathcal{F}_{\text{simp}}] \geq 1 - \delta$$

where $U \in \{0, 1\}^n$ is uniform random and independent of Z . Assume that we can explicitly compute a value p such that for every $i \in [n]$, we have $\mathbb{E}[Z_i] \geq p$. Suppose also that there is an explicit δ -PRG for $\mathcal{F}_{\text{simp}}$ with seed length s' . Then there is an explicit PRG for \mathcal{F} with seed length $t \cdot (s + s')$ and error $O(t\delta)$, where $t = \lceil p^{-1} \ln(n/\delta) \rceil$.

For example, applying the Ajtai-Wigderson framework to Lemma 5.3.2 gives the following PRG for \mathbf{AC}^0 .

Corollary 5.3.4 (PRG for \mathbf{AC}^0 via the Ajtai-Wigderson framework). *For every $n, m, d \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for size- m depth- d \mathbf{AC}^0 circuits with seed length*

$$O(\log m)^{d-1} \cdot \tilde{O}(\log(mn/\varepsilon) \cdot \log(n/\varepsilon)).$$

⁷Note that Lyu actually proved a fully derandomized multi-switching lemma [Lyu22], but we only use the weaker version where the assigned bits are truly random.

When $m = \text{poly}(n)$ and $\varepsilon = 1/\text{poly}(n)$, the seed length in [Corollary 5.3.4](#) is only $O(\log n)^{d+O(1)}$, which is superior to the $O(d)$ exponents in the seed lengths of the PRGs that we saw previously ([Sections 2.6, 4.2 and 5.1](#)). Trevisan and Xue were the first to achieve exponent $d + O(1)$ [[TX13](#)]. Several subsequent papers improved on their seed length [[Tal17](#); [ST19a](#); [Kel21](#); [Lyu22](#)], and the current best seed length is achieved by Lyu [[Lyu22](#)]. Let d be constant and let $m \geq n$. Using a sophisticated variant of the Ajtai-Wigderson framework, Lyu designed an explicit ε -PRG for size- m depth- d \mathbf{AC}^0 circuits with seed length $\tilde{O}(\log^{d-1} m \cdot \log(m/\varepsilon))$, which is quite close to the lack-of-lower-bounds barrier of $\Theta(\log^{d-1} m \cdot \log(1/\varepsilon))$ (see [Section 4.1.2](#)). The remaining gap between the two bounds is most pronounced for the case $d = 2$. The best PRGs for CNFs and DNFs have seed length $\tilde{O}(\log(m/\varepsilon) \cdot \log m)$ [[DETT10](#); [Tal17](#)], whereas the lack-of-lower-bounds barrier allows for a seed length as low as $O(\log m \cdot \log(1/\varepsilon))$.

Open Problem 5.3.5 (Better PRGs for CNFs and DNFs). *Design an explicit PRG for polynomial-size CNFs and DNFs on n variables with error 0.1 and seed length $o(\log^2 n)$.*

In the remainder of this section, we explain the Ajtai-Wigderson framework in its simplest form, i.e., we prove [Theorem 5.3.3](#).

5.3.2 Restrictions that preserve expectation

The first step of the proof of [Theorem 5.3.3](#) is to construct a *fully pseudorandom* restriction that *preserves the expectations* of functions in \mathcal{F} , as defined next.

Definition 5.3.6 (Preserving expectation). *Let $f: \{0, 1\}^n \rightarrow \mathbb{R}$, and let R be a random variable distributed over $\{0, 1, \star\}^n$. We say that R preserves the expectation of f to within ε , or ε -preserves the expectation of f , if*

$$|\mathbb{E}[f|_R(U)] - \mathbb{E}[f]| \leq \varepsilon.$$

Here U is independent of R and distributed uniformly over $\{0, 1\}^n$.

Recall that we are assuming that f simplifies under the partially-pseudorandom restriction $U \circledast Z$. To construct a restriction that preserves the expectation of f , we *replace the truly random bits with stars, and we replace the stars with pseudorandom bits*. To explain further, for $y \in \{0, 1\}^n$, let \bar{y} denote the string obtained by flipping each bit of y , i.e., $\bar{y}_i = 1 - x_i$. Observe that the \circledast operation ([Definition 5.3.1](#)) and the \circ operation ([Definition 5.0.1](#)) satisfy the following “duality” condition:

Fact 5.3.7 (Restriction duality). *For any $x, y, z \in \{0, 1\}^n$, we have $(x \circledast y) \circ z = (z \circledast \bar{y}) \circ x$.*

(See [Figure 5.4](#).) As a consequence of [Fact 5.3.7](#), whenever we have a simplification-under-restriction lemma with a derandomized set of \star positions ([Equation \(5.6\)](#)), there is a related restriction that preserves expectations:

Lemma 5.3.8 (Simplification \implies preserving expectation). *Let $\mathcal{F}_{\text{simp}}$ be a class of functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$. Let X, Z, U be independent random variables taking values in $\{0, 1\}^n$, where U is uniform and X fools $\mathcal{F}_{\text{simp}}$ with error ε . Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$, and assume that*

$$\Pr[f|_{U \circledast Z} \in \mathcal{F}_{\text{simp}}] \geq 1 - \delta.$$

Then $X \circledast \bar{Z}$ preserves the expectation of f to within $\varepsilon + \delta$.




Figure 5.4: To compute $(x \star y) \circ z$, we use y to partition the coordinates into two parts. The coordinates in one part get their values from x , while the coordinates in the other part get their values from z . Thus, if we swap the roles of x and z and flip each bit of y , nothing changes. Here we depict the case $y = 01011001$.

Proof. We must show that $(X \star \bar{Z}) \circ U$ fools f . By [Fact 5.3.7](#),

$$\begin{aligned} f((X \star \bar{Z}) \circ U) &= f((U \star Z) \circ X) \\ &= f|_{U \star Z}(X). \end{aligned}$$

By assumption, except with probability δ , $f|_{U \star Z} \in \mathcal{F}_{\text{simp}}$, and in this case, X fools the restricted function with error ε . \square

Observe that if $U \star Z$ (the restriction that causes simplification) has \star -probability p , then $X \star \bar{Z}$ (the restriction that preserves expectations) has \star -probability $1 - p$. When we are trying to prove simplification under restrictions, we want the \star -probability to be as large as possible, whereas when we are trying to prove preservation of expectation, we want the \star -probability to be as small as possible.

5.3.3 Iterated restrictions

[Lemma 5.3.2](#) provides us with a fully pseudorandom restriction R that preserves the expectation of $f \in \mathcal{F}$. Because R is fully pseudorandom, we can afford to sample it and apply it, so f becomes $f|_R$. The next lemma says that if we then sample another copy of R and further restrict the restricted function, we continue preserving its expectation.

Lemma 5.3.9 (Composing restrictions that preserve expectations). *Let \mathcal{F} be a class of functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$ that is closed under restriction. Let $R^{(1)}, \dots, R^{(t)}$ be independent random variables distributed over $\{0, 1, \star\}^n$, assume that each $R^{(i)}$ preserves the expectation of every $f \in \mathcal{F}$ to within ε , and let $R = R^{(1)} \circ \dots \circ R^{(t)}$. Then R preserves the expectation of every $f \in \mathcal{F}$ to within $\varepsilon \cdot t$.*

Proof. Let us prove it by induction on t . The base case $t = 1$ is trivial. For the inductive step, let $R^{(<t)} = R^{(1)} \circ \dots \circ R^{(t-1)}$ and assume that $R^{(<t)}$ preserves the expectation of every $f \in \mathcal{F}$ to within $\varepsilon \cdot (t - 1)$. Fix some $f \in \mathcal{F}$, and let $F = f|_{R^{(<t)}}$. Since \mathcal{F} is closed under restriction, we have $F \in \mathcal{F}$

with probability 1. Sample $U \in \{0, 1\}^n$ uniformly at random. Then

$$\begin{aligned}
|\mathbb{E}[f|_R(U)] - \mathbb{E}[f]| &= |\mathbb{E}[f(R^{(<t)} \circ R^{(t)} \circ U)] - \mathbb{E}[f]| \\
&= |\mathbb{E}[F|_{R^{(t)}}(U)] - \mathbb{E}[f]| \\
&\leq |\mathbb{E}[F|_{R^{(t)}}(U)] - \mathbb{E}[F(U)]| + |\mathbb{E}[F(U)] - \mathbb{E}[f]| \\
&\leq \mathbb{E}_F \left[\left| \mathbb{E}_{R^{(t)}, U} [F|_{R^{(t)}}(U)] - \mathbb{E}_U [F(U)] \right| \right] + |\mathbb{E}[f|_{R^{(<t)}}(U)] - \mathbb{E}[f]| \\
&\leq \mathbb{E}_F[\varepsilon] + \varepsilon \cdot (t - 1) = \varepsilon \cdot t. \quad \square
\end{aligned}$$

To get a full PRG, we will take t to be large enough that with high probability, the composed restriction R in [Lemma 5.3.9](#) assigns values to *all* variables.

Lemma 5.3.10 (Preserving expectation \implies PRG). *Let \mathcal{F} be a class of functions $f: \{0, 1\}^n \rightarrow \mathbb{R}$ that is closed under restrictions. Suppose there is a distribution R over $\{0, 1, \star\}^n$ such that R can be explicitly sampled using s truly random bits, and R preserves the expectation of every $f \in \mathcal{F}$ to within δ , and for every coordinate i , $\Pr[R_i = \star] \leq 1 - p$. Then there is an explicit PRG for \mathcal{F} with seed length $s \cdot t$ and error $O(t\delta)$, where $t = \lceil p^{-1} \ln(n/\delta) \rceil$.*

Proof. Let $R^{\circ t} = R^{(1)} \circ \dots \circ R^{(t)}$, where $R^{(1)}, \dots, R^{(t)}$ are independent copies of R . Our PRG outputs the string $R^{\circ t} \circ 0^n$. This PRG clearly has seed length $s \cdot t$. By [Lemma 5.3.9](#), the restriction $R^{\circ t}$ preserves the expectation of every $f \in \mathcal{F}$ to within error $t \cdot \delta$. Furthermore, the probability that there are any stars remaining in $R^{\circ t}$ is bounded by

$$\Pr[R^{\circ t} \notin \{0, 1\}^n] \leq \sum_{i=1}^n \Pr[R_i^{\circ t} = \star] = \sum_{i=1}^n \Pr[R_i = \star]^t \leq \sum_{i=1}^n (1 - p)^t \leq \delta.$$

Consequently, $R^{\circ t} \circ 0^n$ fools \mathcal{F} with error $(t + 1) \cdot \delta$. □

[Theorem 5.3.3](#) (the Ajtai-Wigderson reduction) follows from [Lemmas 5.3.8](#) and [5.3.10](#).

5.4 The Forbes-Kelley generator for ROBPs

In [Section 5.2](#), we presented a PRG for arbitrary-order *permutation* ROBPs. In this section, we present Forbes and Kelley's PRGs [\[FK18\]](#), which fool *general* arbitrary-order ROBPs (without the permutation assumption). The PRGs use a generalization of the Ajtai-Wigderson framework ([Section 5.3](#)). We begin by explaining the generalization of the framework.

5.4.1 Pseudorandomness plus noise

Let \mathcal{F} be a class of functions that we wish to fool. Recall the first step of the Ajtai-Wigderson framework: we showed ([Lemma 5.3.8](#)) that if functions in \mathcal{F} simplify with high probability under partially-pseudorandom restrictions, then we get a *fully pseudorandom* restriction that *preserves the expectation* of each $f \in \mathcal{F}$. We now refine that analysis to get an "if and only if" condition (actually three equivalent conditions).

Lemma 5.4.1 (Characterizing preservation of expectation). *Let X, Z, U be mutually independent random variables, where each is distributed over $\{0, 1\}^n$ and U is uniform random. Let $f: \{0, 1\}^n \rightarrow \mathbb{R}$ be a function and let $\varepsilon > 0$. The following are equivalent.*

1. (*Preservation of Expectation*) The restriction $X \star Z$ preserves the expectation of f to within error ε , i.e.,

$$|\mathbb{E}[f|_{X \star Z}(U)] - \mathbb{E}[f]| \leq \varepsilon.$$

2. (*Simplification on Average*) The distribution X fools g with error ε , where

$$g(x) \stackrel{\text{def}}{=} \mathbb{E}_{Z,U}[f|_{U \star \bar{Z}}(x)]. \quad (5.7)$$

3. (*Pseudorandomness Plus Noise*) The distribution $X + (Z \wedge U)$ fools f with error ε , i.e.,

$$|\mathbb{E}[f(X + (Z \wedge U))] - \mathbb{E}[f]| \leq \varepsilon, \quad (5.8)$$

where $+$ denotes addition over \mathbb{F}_2^n and \wedge denotes coordinatewise conjunction.

Proof. (1 \iff 2) By [Fact 5.3.7](#), we have

$$f|_{X \star Z}(U) = f((U \star \bar{Z}) \circ X) = f|_{U \star \bar{Z}}(X).$$

Therefore, $\mathbb{E}[f|_{X \star Z}(U)] = \mathbb{E}[g(X)]$. Furthermore, because U is uniform random, we have $\mathbb{E}[g] = \mathbb{E}[f]$.

(2 \iff 3) Because U is uniform random, the random variables $X + (Z \wedge U)$ and $(U \star \bar{Z}) \circ X$ are identically distributed. (In each case, we start with X and then we randomize the bits where $Z_i = 1$.) Therefore, $\mathbb{E}[f(X + (Z \wedge U))] = \mathbb{E}[g(X)]$. As mentioned previously, the fact that U is uniform random also implies that $\mathbb{E}[f] = \mathbb{E}[g]$. \square

Recall that the *noise operator* with parameter p is defined by $T_p f(x) = \mathbb{E}[f|_{\mathcal{R}_p}(x)]$. The function g defined in [Equation \(5.7\)](#) can be understood as the result of applying a *partially derandomized noise operator* to f . Thus, [Item 2](#) says that f *simplifies on average under partially-pseudorandom restrictions*. This condition is a combination of what we studied in [Section 5.2](#) (simplification on average) and what we studied in [Section 5.3](#) (simplification under partially-pseudorandom restrictions). A form of the perspective embodied by [Item 2](#) was first studied by Gopalan, Meka, Reingold, Trevisan, and Vadhan [[GMRTV12](#)].

In this section, it will not be so useful to think in terms of “simplification.” Instead, it will be more productive to reason directly about *fooling f using a partially-pseudorandom distribution*, along the lines [Item 3](#). Intuitively, establishing [Equation \(5.8\)](#) is easier than trying to design a PRG in one shot, because the helpful “noise vector” $Z \wedge U$ contributes a considerable amount of true randomness into the picture. This perspective originates in work by Haramaty, Lee, and Viola [[HLV18](#)]. We adopt this perspective for the rest of this section.

Once we obtain random variables $X, Z \in \mathbb{F}_2^n$ satisfying [Equation \(5.8\)](#), we can repeat the same procedure iteratively to get a PRG for \mathcal{F} . This is because for any fixed $x, z \in \mathbb{F}_2^n$, fooling the new function $g(y) = f(x + (z \wedge y))$ is equivalent to fooling $f|_{x \star z}$ (see [Lemma 5.3.10](#)). We will explain an important example of this approach next, namely the Forbes-Kelley PRG for arbitrary-order ROBPs.

5.4.2 A Fourier decomposition lemma for ROBPs

We return now to the model of *arbitrary-order ROBPs* introduced in [Section 5.2](#). Unlike in [Section 5.2](#), in this section we will not assume that the ROBP is regular. From the Fourier growth bound [Theorem 5.2.14](#), it follows that one can fool constant-width arbitrary-order ROBPs with error

$1/\text{poly}(n)$ and seed length $\text{polylog}(n)$. Forbes and Kelley used the iterated restrictions framework to show that we can fool even *polynomial-width arbitrary-order ROBPs* with error $1/\text{poly}(n)$ and seed length $O(\log^3 n)$ [FK18]. This implies the first PRG with polylogarithmic seed length for *read-once formulas with constant fan-in* over an arbitrary basis, since every such formula can be computed by a polynomial-width ROBP under some input order [BPW11].

The starting point for Forbes and Kelley's work is the following Fourier decomposition lemma for ROBPs. For simplicity, we assume that the ROBP uses the standard variable ordering (note that permuting variables just permutes Fourier coefficients in the obvious way).

Lemma 5.4.2 (Forbes-Kelley decomposition of ROBPs). *Let f be a length- n width- w standard-order ROBP with layers V_0, V_1, \dots, V_n . Then*

$$f(x) = L(x) + H(x),$$

where

$$L(x) = \sum_{S:|S|<k} \widehat{f}(S)\chi_S(x),$$

and

$$H(x) = \sum_{i=1}^n \sum_{v \in V_i} H_v(x) f_{v \rightarrow}(x),$$

where $H_v(x) = \sum_{S \subseteq [i]:|S|=k, i \in S} \widehat{f}_{\rightarrow v}(S)\chi_S(x)$.

Proof. Since $L(x)$ is exactly the degree $< k$ part of the Fourier expansion of f , we just need to show that

$$H(x) = \sum_{S \subseteq [n]:|S| \geq k} \widehat{f}(S)\chi_S(x).$$

For every S with $|S| \geq k$, let $i(S)$ denote the k -th smallest index that appears in S , let $S_L := S \cap [i(S)]$, and let $S_R := S \setminus S_L$. Note that $|S_L| = k$ and $i(S) \in S_L$. We have

$$\widehat{f}(S) = \sum_{v \in V_{i(S)}} \widehat{f}_{\rightarrow v}(S_L) \cdot \widehat{f}_{v \rightarrow}(S_R),$$

where we used the fact that $f(x) = \sum_{i \in V_{i(S)}} f_{\rightarrow v}(x) \cdot f_{v \rightarrow}(x)$. Thus,

$$\begin{aligned} \sum_{S \subseteq [n]:|S| \geq k} \widehat{f}(S)\chi_S(x) &= \sum_{i=1}^n \sum_{S:i(S)=i} \widehat{f}(S)\chi_S(x) \\ &= \sum_{i=1}^n \sum_{S:i(S)=i} \sum_{v \in V_i} \widehat{f}_{\rightarrow v}(S_L) \cdot \widehat{f}_{v \rightarrow}(S_R)\chi_{S_L}(x)\chi_{S_R}(x) \\ &= \sum_{i=1}^n \sum_{v \in V_i} \sum_{\substack{S_L \subseteq [i] \\ |S_L|=k \\ i \in S_L}} \sum_{S_R \subseteq [n] \setminus [i]} \widehat{f}_{\rightarrow v}(S_L) \cdot \widehat{f}_{v \rightarrow}(S_R)\chi_{S_L}(x)\chi_{S_R}(x) \\ &= \sum_{i=1}^n \sum_{v \in V_i} H_v(x) \sum_{S_R \subseteq [n] \setminus [i]} \widehat{f}_{v \rightarrow}(S_R)\chi_{S_R}(x) \\ &= \sum_{i=1}^n \sum_{v \in V_i} H_v(x) f_{v \rightarrow}(x) \\ &= H(x). \end{aligned} \quad \square$$

5.4.3 Pseudorandom restrictions that preserve the expectation of ROBPs

The key point of the construction will be to analyze ROBPs under bounded independent restrictions. In the proposition below, once again we assume for simplicity that the ROBP reads its variables in the standard order; this is without loss of generality because k -wise uniformity is preserved under variable permutations.

Proposition 5.4.3 (Preserving the expectation of ROBPs). *Suppose $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is computed by a width- w standard-order ROBP. Suppose that X, Z, U are independent random variables, where X is $2k$ -wise uniform, Z is k -wise uniform, and U is uniform. Then*

$$\left| \mathbb{E}_{X,Z,U}[f(X + Z \wedge U)] - \mathbb{E}[f] \right| \leq \frac{wn}{2^{k/2}}.$$

Remark 5.4.4. *For our current project of designing PRGs for ROBPs, the fact that Z is pseudorandom in Proposition 5.4.3 is not crucial. After all, if we take Z to be truly random, then Proposition 5.4.3 says that X fools $T_{1/2}f$, where $T_{1/2}$ is the noise operator with parameter $p = 1/2$. From there, the polarizing random walks framework gives a PRG (see Sections 5.1 and 5.2). Since p is a constant, the seed length from the polarizing random walks framework is the same as the seed length from the iterated restrictions framework for this case. (In general, the polarizing walks framework has a worse dependence on p .)*

That being said, there are other benefits to the iterated restrictions paradigm besides seed length. In particular, the fact that Z is pseudorandom will be crucial in Section 5.5, which builds on this section using an “early termination” approach.

Proof. The key is to utilize Lemma 5.4.2. We have

$$\begin{aligned} & \left| \mathbb{E}_{X,Z,U}[f(X + Z \wedge U)] - \mathbb{E}[f] \right| \\ &= \left| \mathbb{E}_{X,Z,U}[L(X + Z \wedge U) + H(X + Z \wedge U)] - \widehat{f}(\emptyset) \right| \\ &= \left| \mathbb{E}_{X,Z,U}[H(X + Z \wedge U)] \right|, \end{aligned}$$

where the last equality uses the fact that

$$\mathbb{E}_{X,Z,U}[L(X + Z \wedge U)] = \widehat{f}(\emptyset),$$

since $X + Z \wedge U$ is k -wise uniform. It is left to bound $|\mathbb{E}_{X,Z,U}[H(X + Z \wedge U)]|$. Note that,

$$\begin{aligned} & \left| \mathbb{E}_{X,Z,U}[H(X + Z \wedge U)] \right| \\ &= \left| \sum_{i=1}^n \sum_{v \in V_i} \mathbb{E}_{X,Z,U}[H_v(X + Z \wedge U) f_{v \rightarrow}(X + Z \wedge U)] \right| \\ &= \left| \sum_{i=1}^n \sum_{v \in V_i} \mathbb{E}_{X,Z} \left[\mathbb{E}_U[H_v(X + Z \wedge U)] \mathbb{E}_U[f_{v \rightarrow}(X + Z \wedge U)] \right] \right| \\ &\leq \sum_{i=1}^n \sum_{v \in V_i} \mathbb{E}_{D,T} \left[\left| \mathbb{E}_U[H_v(X + Z \wedge U)] \right| \right] \\ &\leq nw \cdot \max_v \mathbb{E}_{X,Z} \left[\left| \mathbb{E}_U[H_v(X + Z \wedge U)] \right| \right], \end{aligned}$$

where the first inequality uses the triangle inequality and the fact that $|\mathbb{E}_U[f_{v \rightarrow}(X + Z \wedge U)]| \leq 1$. It is thus sufficient to bound $\mathbb{E}_{X,Z} [|\mathbb{E}_U[H_v(X + Z \wedge U)]|]$ for every i and $v \in V_i$.

$$\begin{aligned}
& \mathbb{E}_{X,Z} \left[\left| \mathbb{E}_U[H_v(X + Z \wedge U)] \right|^2 \right] \\
& \leq E_{X,Z} \left[\left(\mathbb{E}_U[H_v(X + Z \wedge U)] \right)^2 \right] \\
& = \mathbb{E}_{X,Z,U,U'} \left[\sum_{\substack{S,S' \subseteq [i]: \\ i \in S,S', |S|=|S'|=k}} \widehat{f}_{\rightarrow v}(S) \widehat{f}_{\rightarrow v}(S') \chi_S(X + Z \wedge U) \chi_{S'}(X + Z \wedge U') \right] \\
& = \sum_{\substack{S,S' \subseteq [i]: \\ i \in S,S', |S|=|S'|=k}} \widehat{f}_{\rightarrow v}(S) \widehat{f}_{\rightarrow v}(S') \mathbb{E}_{X,Z} \left[\mathbb{E}_U[\chi_S(X + Z \wedge U)] \mathbb{E}_{U'}[\chi_{S'}(X + Z \wedge U')] \right]
\end{aligned}$$

where the first step is the Cauchy-Schwarz inequality. Now note that, whenever Z has a 1 coordinate in S or S' , then $\mathbb{E}_U[\chi_S(X + Z \wedge U)] \mathbb{E}_{U'}[\chi_{S'}(X + Z \wedge U')] = 0$. Otherwise when Z is all zeros on both S and S' coordinates, we have $\chi_S(X + Z \wedge U) = \chi_S(X)$ and $\chi_{S'}(X + Z \wedge U') = \chi_{S'}(X)$. Now since X is $2k$ -wise uniform, in this case we get

$$\mathbb{E}_X \left[\mathbb{E}_U[\chi_S(X + Z \wedge U)] \mathbb{E}_{U'}[\chi_{S'}(X + Z \wedge U')] \right] = \mathbb{E}_X [\chi_S(X) \cdot \chi_{S'}(X)] = \begin{cases} 1 & \text{if } S = S' \\ 0 & \text{if } S \neq S'. \end{cases}$$

Putting these facts together, we get

$$\begin{aligned}
& \sum_{\substack{S,S' \subseteq [i]: \\ i \in S,S', |S|=|S'|=k}} \widehat{f}_{\rightarrow v}(S) \widehat{f}_{\rightarrow v}(S') \mathbb{E}_{X,Z} \left[\mathbb{E}_U[\chi_S(X + Z \wedge U)] \mathbb{E}_{U'}[\chi_{S'}(X + Z \wedge U')] \right] \\
& = \sum_{\substack{S \subseteq [i]: \\ i \in S, |S|=k}} \widehat{f}_{\rightarrow v}(S)^2 \Pr_Z[Z_i = 0 \forall i \in S] \\
& = 2^{-k} \cdot \sum_{\substack{S \subseteq [i]: \\ i \in S, |S|=k}} \widehat{f}_{\rightarrow v}(S)^2 \leq 2^{-k},
\end{aligned}$$

where the second equality is due to T being k -wise uniform, and the last inequality follows from Parseval's identity. \square

Given [Proposition 5.4.3](#), we can get a full ε -PRG for arbitrary-order ROBPs by choosing $k = O(\log(wn/\varepsilon))$ and applying the generic reduction [Lemma 5.3.10](#). Using efficient constructions of k -wise and $2k$ -wise uniform distributions, an individual restriction can be sampled explicitly using $O(k \log n) = O(\log n \cdot \log(\frac{nw}{\varepsilon}))$ truly random bits. Therefore, the overall seed length is as follows.

Theorem 5.4.5 (PRGs for arbitrary-order ROBPs [[FK18](#)]). *For every $n, w \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for width- w length- n arbitrary-order ROBPs with seed length $O(\log n \cdot \log(n/\varepsilon) \cdot \log(nw/\varepsilon))$. In particular, when $w = \text{poly}(n)$ and $\varepsilon = 1/\text{poly}(n)$, the seed length is $O(\log^3 n)$.*

For small values of ε , a better seed length of $O(\log(nw/\varepsilon) \cdot \log^2 n)$ can be achieved by terminating the iterative restriction process after $O(\log n)$ restrictions instead of $O(\log(n/\varepsilon))$ restrictions and observing that the restricted function is an $O(\log(nw/\varepsilon))$ -junta with high probability. See Forbes and Kelley's work for details [[FK18](#)].

5.4.4 A better generator for the small-width setting

Forbes and Kelley showed how to achieve a better seed length when w is small. The construction is similar, except that k -wise uniform distributions are replaced by small-bias distributions.⁸

Proposition 5.4.6 (Preserving the expectation of ROBPs that have low level- k Fourier weight [FK18, Lemma 7.2]). *Suppose $f: \{0, 1\}^n \rightarrow \{0, 1\}$ is computed by a width- w standard-order ROBP. Moreover, suppose that X and Z are independent β -biased random variables distributed over $\{0, 1\}^n$. Let $k \in \mathbb{N}$, and let $L = \sum_{|S|=k} |\widehat{f}(S)|$. Then the restriction $X \star Z$ preserves the expectation of f to within*

$$O\left(\left(2^{-k/2} + \sqrt{\beta} \cdot (L + 2^{k/4})\right) \cdot n \cdot w\right).$$

The proof of Proposition 5.4.6 is similar to the proof of Proposition 5.4.3, and we omit it. The point is that Proposition 5.4.6 allows us to take advantage of bounds on the Fourier growth of f . Plugging Theorem 5.2.14 into Proposition 5.4.6 and choosing $k = \Theta(\log(wn/\varepsilon))$ yields the following.

Proposition 5.4.7 (Preserving the expectation of low-width ROBPs). *For every $w, n \in \mathbb{N}$ and $\varepsilon > 0$, there is a value*

$$\beta = 2^{-O(\log(wn/\varepsilon) \cdot w \cdot \log \log n)}$$

such that if X and Z are independent β -biased random variables distributed over $\{0, 1\}^n$, then the restriction $X \star Z$ preserves the expectation of any width- w ROBP $f: \{0, 1\}^n \rightarrow \{0, 1\}$ to within ε .

Note that the restriction $X \star Z$ has \star -probability $1/2$, so we are assigning values to roughly half the input bits. Furthermore, when w is a constant, the restriction $X \star Z$ can be sampled using $\tilde{O}(\log(n/\varepsilon))$ truly random bits. Iterating for $O(\log n)$ steps like before leads to the following PRG.

Theorem 5.4.8 ([FK18]). *There is an explicit PRG for width- w length- n arbitrary-order ROBPs with error ε and seed length $O(w \cdot \log(nw/\varepsilon) \cdot \log n \cdot \log \log n)$.*

5.5 PRGs for read-once CNFs via early termination

Let \mathcal{F} be a class of functions that we wish to fool. In Sections 5.3 and 5.4, our approach for fooling \mathcal{F} was to first design a pseudorandom restriction R that approximately preserves the expectations of functions in \mathcal{F} . Then, we iteratively applied many copies of this restriction, assigning values to more and more variables. If R assigns values to a p -fraction of the variables, then we perform roughly $p^{-1} \cdot \log n$ many iterations. Thus, if each copy of R costs s truly random bits, then this approach leads to a final seed length of roughly $p^{-1} \cdot s \cdot \log n$.

In this section, we develop a refined version of the iterated restrictions paradigm that can lead to an improved seed length in some cases. The idea is, we start by composing t copies of R , say $R' = R^{(1)} \circ \dots \circ R^{(t)}$, where the number of iterations (t) is significantly smaller than $p^{-1} \cdot \log n$, such as perhaps $t = p^{-1} \cdot \log \log n$. Then, we prove that functions in \mathcal{F} simplify under this composed restriction R' . That is, we show that

$$\forall f \in \mathcal{F}, \Pr[f|_{R'} \in \mathcal{F}_{\text{simp}}] \geq 1 - \delta, \tag{5.9}$$

where $\mathcal{F}_{\text{simp}}$ is some class of “simpler” functions. Consequently, there is no need to continue sampling copies of R . Instead, we can use a PRG for $\mathcal{F}_{\text{simp}}$ to finish the job, taking advantage of

⁸In Forbes and Kelley’s work [FK18], they take the star-set T to be almost k -wise uniform; this is implied by the small-bias condition (see Theorem 2.3.2.)

the “simplicity” of the restricted function $f|_{R'}$. This leads to a final seed length of $s \cdot t + s'$, where s' is the seed length for fooling $\mathcal{F}_{\text{simp}}$.

Observe that the restriction R' in this framework is fully pseudorandom. Thus, the key challenge of this approach is that it requires proving a *fully-derandomized* simplification-under-restrictions lemma (Equation (5.9)). In Section 5.5.1, we show an example of a fully-derandomized simplification-under-restriction lemma, for the class of read-once CNF formulas. Then, in Section 5.5.2, we explain how to use that lemma to design a near-optimal PRG for such formulas.

5.5.1 Simplification of read-once CNFs under fully-pseudorandom restrictions

Recall that a CNF is a conjunction of *clauses*, each of which is a disjunction of literals. We will show that under a suitable pseudorandom restriction, a read-once CNF is likely to simplify, in the sense that the restricted formula only has a few remaining clauses. We begin with the following convenient definition, which generalizes the notion of γ -almost k -wise uniformity.

Definition 5.5.1 (*k -wise γ -close distributions*). *Let Σ be an alphabet, let $n \in \mathbb{N}$, and let X, Y be distributions over Σ^n . Let $k \in \mathbb{N}$ and $\gamma > 0$. We say that X is k -wise γ -close to Y if for every $S \subseteq [n]$ with $|S| \leq k$, the substrings X_S and Y_S have total variation distance at most γ .*

Recall that the *width* of a clause is the number of literals in the clause. A width- w CNF is a CNF in which each clause has width at most w . We begin by analyzing the effect of a “mild” pseudorandom restriction, i.e., a restriction that only assigns values to a constant fraction of the input variables. Intuitively, under such a restriction, the width of a read-once CNF should decrease by a constant factor, because in any given clause, a constant fraction of the variables should be assigned values. We show that this indeed occurs in all but a few “exceptional” clauses.

Lemma 5.5.2 (Simplification of read-once CNFs under mild fully-pseudorandom restrictions). *For every $w \in \mathbb{N}$ and $\delta > 0$, there is a value $k = O(w + \log(1/\delta))$ such that the following holds. Let f be a width- w read-once CNF. Let $\gamma = 2^{-4k}$, let $p \leq 2^{-10}$, and let \tilde{R} be a distribution over $\{0, 1, \star\}^n$ that is k -wise γ -close to \mathcal{R}_p . Then with probability $1 - \delta$, the restricted function $f|_{\tilde{R}}$ can be computed by a read-once CNF of the form $f|_{\tilde{R}} = f_{\text{narrow}} \wedge f_{\text{small}}$, where f_{narrow} is a read-once CNF of width at most $w/2$ and f_{small} is a read-once CNF with at most $O(\log^2(1/\delta))$ clauses.*

To prove Lemma 5.5.2, we rely on the following tail bound, which we cite without proof.

Theorem 5.5.3 (Tail bound for almost k -wise independent random variables [DMRTV21]). *Let $q > 0$ and let $X \in \{0, 1\}^m$, where X_1, \dots, X_m are independent random variables with $\mathbb{E}[X_i] \geq q$ for each i . Let k be an even positive integer, let $\gamma > 0$, and let \tilde{X} be k -wise γ -close to X . Then⁹*

$$\Pr \left[\tilde{X}_1 = \tilde{X}_2 = \dots = \tilde{X}_m = 0 \right] \leq \left(\frac{16k}{qm} \right)^{k/2} + 2k \cdot \gamma \cdot \left(\frac{1}{q} \right)^k.$$

The specific bound of Theorem 5.5.3 appears in work by Doron, Meka, Reingold, Tal, and Vadhan [DMRTV21] as a corollary of earlier work by Steinke, Vadhan, and Wan [SVW17]; a similar bound appears in work by Celis, Reingold, Segev, and Wieder [CRSW13]. Now we prove Lemma 5.5.2 using Theorem 5.5.3.

⁹The statement in Doron, Meka, Reingold, Tal, and Vadhan’s work [DMRTV21] includes an extra assumption $k \leq qm$. This assumption is not necessary, because if $k > qm$, then the conclusion of the theorem is trivial.

Proof of Lemma 5.5.2. Let k be a multiple of w such that $32k \cdot 2^{-k/2} \leq \delta$. Let m be the number of clauses in f , say $f = C_1 \wedge \cdots \wedge C_m$. We consider three cases based on the value of m . For the first case, suppose that $m \leq 8^w$. For $j \in [m]$, let X_j be the indicator for the bad event that there are more than $w/2$ variables that are read by C_j and kept alive by R . Let f_{small} consist of all the clauses with $X_j = 1$ and let f_{narrow} consist of all the clauses with $X_j = 0$; our job is to show that f_{small} has few clauses with high probability. Let w_j be the number of variables that C_j reads, so $w_j \leq w$. If \tilde{R} is a *truly* random restriction $\tilde{R} = \mathcal{R}_p$, then

$$\mathbb{E}[X_j] \leq \binom{w_j}{w/2} \cdot p^{w/2} \leq 2^w \cdot p^{w/2} \leq 2^{-4w}.$$

For any set $S \subseteq [m]$ of size $|S| = k/w$, the clauses $\{C_j\}_{j \in S}$ read a total of at most k variables, so when \tilde{R} is a *pseudorandom* restriction (namely k -wise γ -close to \mathcal{R}_p), we have

$$\Pr[X_S = 1^S] \leq \gamma + (2^{-4w})^{k/w} = 2 \cdot 2^{-4k}.$$

Therefore, by the union bound,

$$\Pr \left[\sum_j X_j \geq k/w \right] \leq \binom{m}{k/w} \cdot 2 \cdot 2^{-4k} \leq 8^{w \cdot k/w} \cdot 2 \cdot 2^{-4k} = 2 \cdot 2^{-k} \leq \delta/(16k).$$

Note that $k/w = O(1 + \log(1/\delta)/w) \leq O(\log(1/\delta))$, so indeed, with high probability, f_{small} has at most $O(\log(1/\delta))$ clauses.

Next, for the second case, suppose that $8^w < m \leq 8^w \cdot 16(k/w)$. In this case, write $f = f_1 \wedge \cdots \wedge f_{16(k/w)}$, where each f_j is a read-once CNF of width at most w with at most 8^w clauses. We apply the previous argument to each f_j . By the union bound, with probability $1 - \delta$, each $f_j|_{\tilde{R}}$ can be written as $f_{\text{narrow}}^{(j)} \wedge f_{\text{small}}^{(j)}$, where $f_{\text{narrow}}^{(j)}$ has width at most $w/2$ and $f_{\text{small}}^{(j)}$ has at most $O(\log(1/\delta))$ clauses. Define $f_{\text{narrow}} = \bigwedge_j f_{\text{narrow}}^{(j)}$ and $f_{\text{small}} = \bigwedge_j f_{\text{small}}^{(j)}$, and observe that f_{small} has at most $O(\log(1/\delta) \cdot k/w) \leq O(\log^2(1/\delta))$ clauses.

For the final case, suppose that $m > 8^w \cdot 16(k/w)$. In this case, we will show that with high probability, the restricted function is identically zero. Let $X = (X_1, \dots, X_m)$, where X_j indicates whether the clause C_j is falsified by a truly random restriction, i.e., $X_j = 1 \iff C_j|_{\mathcal{R}_p} \equiv 0$. Similarly, let $\tilde{X} = (\tilde{X}_1, \dots, \tilde{X}_m)$, where \tilde{X}_j indicates whether $C_j|_{\tilde{R}} \equiv 0$. Observe that X_1, \dots, X_m are independent and $\mathbb{E}[X_j] \geq (\frac{1-p}{2})^w \geq 4^{-w}$. Meanwhile, the vector \tilde{X} is (k/w) -wise γ -close to X . Therefore, by [Theorem 5.5.3](#),

$$\begin{aligned} \Pr[f|_R \neq 0] &= \Pr[\tilde{X} = 0^m] \leq \left(\frac{16k/w}{4^{-w}m} \right)^{k/(2w)} + 2(k/w) \cdot \gamma \cdot 4^{w(k/w)} \\ &\leq 2^{-k/2} + 2(k/w) \cdot 2^{-2k} \\ &\leq \delta. \end{aligned} \quad \square$$

Next, we analyze the effect of a pseudorandom restriction with a relatively small \star -probability (more similar to the traditional “switching lemma” regime). By applying [Lemma 5.5.2](#) several times, we show that the number of clauses in a read-once CNF drastically decreases under such a restriction.

Corollary 5.5.4 (Simplification of read-once CNFs under fully-pseudorandom restrictions). *For every $w \in \mathbb{N}$ and $\delta > 0$, there is a value $k = O(w + \log(1/\delta))$ such that the following holds. Let*

f be a width- w read-once CNF. Let $\gamma = 2^{-4k}$, let $p \leq 2^{-10}$, let $t = \lceil \log w \rceil$, and let $\tilde{R}^{(1)}, \dots, \tilde{R}^{(t)}$ be independent random variables, where each $\tilde{R}^{(j)}$ is k -wise γ -close to \mathcal{R}_p . Then with probability $1 - \delta$, the restricted function $f|_{\tilde{R}^{(1)} \circ \dots \circ \tilde{R}^{(t)}}$ can be computed by a read-once CNF with at most $\tilde{O}(\log w \cdot \log^2(1/\delta))$ many clauses.

Proof sketch. We repeatedly apply [Lemma 5.5.2](#) with failure probability δ/t . The first time, we apply it to f , which with high probability becomes $f_{\text{narrow}} \wedge f_{\text{small}}$. The second time, we apply it to f_{narrow} , which becomes $f'_{\text{narrow}} \wedge f'_{\text{small}}$. The third time, we apply it to f'_{narrow} , etc. After t iterations, the width of the “narrow” part drops below 1, so only the “small” parts remain. All together, the “small” parts have $O(t \cdot \log^2(t/\delta))$ many clauses. \square

5.5.2 Iterated restrictions with early termination

So far, we have shown that read-once CNFs simplify under fully-pseudorandom restrictions ([Corollary 5.5.4](#)). Next, we shall use [Corollary 5.5.4](#) to design a PRG for read-once CNFs. Previously, we showed that δ -biased generators ε -fool read-once \mathbf{AC}^0 formulas ([Section 2.5.3](#)); this led to a seed length of $O(\log n \cdot \log(1/\varepsilon))$ for the depth-two case. Now we will show how to achieve near-optimal seed length $\tilde{O}(\log(n/\varepsilon))$.

Theorem 5.5.5 (Near-optimal PRG for read-once depth-two formulas). *For every $n \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for read-once CNFs and DNFs with seed length $\tilde{O}(\log(n/\varepsilon))$.*

[Theorem 5.5.5](#) was first proven by Gopalan, Meka, Reingold, Trevisan, and Vadhan [[GMRTV12](#)]. (Compared to Gopalan, Meka, Reingold, Trevisan, and Vadhan’s original proof [[GMRTV12](#)], the analysis we present here is more similar to later work that considers more general models [[Lee19](#); [DMRTV21](#)].) Later, Doron, Hatami, and Hoza achieved a slightly better seed length of $O(\log n) + \tilde{O}(\log(1/\varepsilon))$ [[DHH20](#)]. It remains an open problem to achieve the optimal seed length for this basic and fundamental class.

Open Problem 5.5.6 (Optimal PRGs for read-once depth-two formulas). *Construct an explicit PRG for read-once CNFs with seed length $O(\log(n/\varepsilon))$.*

As outlined at the beginning of this section, we prove [Theorem 5.5.5](#) using the iterated restrictions paradigm. As the first (and main) step, let us aim for seed length $\tilde{O}(w + \log(n/\varepsilon))$, where w is the width of the formula.

Lemma 5.5.7 (PRGs for bounded-width read-once CNFs). *For every $w, n \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit ε -PRG for width- w read-once CNFs on n input variables with seed length $\tilde{O}(w + \log(n/\varepsilon))$.*

Proof. The PRG outputs $R^{(1)} \circ \dots \circ R^{(t)} \circ Z$, where each $R^{(i)}$ is distributed according to $X \star Y$, X and Y are β -biased, and Z is α -biased for suitable values

$$\begin{aligned} t &= O(\log w) \\ \beta &= 2^{-\Theta(w + \log(n/\varepsilon) \cdot \log \log n)} \\ \alpha &= 2^{-\Theta(\log(1/\varepsilon) \cdot \log \log(w/\varepsilon))}. \end{aligned}$$

Observe that the seed length is $O(t \log(n/\beta) + \log(1/\alpha))$, which is indeed $\tilde{O}(w + \log(n/\varepsilon))$. Our remaining job is to prove correctness. Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be a width- w read-once CNF, and let R be the composed restriction $R = R^{(1)} \circ \dots \circ R^{(t)}$.

The first step of the correctness proof is to show that each individual restriction $R^{(i)}$ preserves the expectation of read-once CNFs. This is a straightforward consequence of Forbes and Kelley’s

work (see [Section 5.4](#)). Indeed, read-once CNFs can be simulated by width-3 arbitrary-order ROBPs. Therefore, by [Proposition 5.4.7](#), each individual restriction $R^{(i)}$ preserves the expectation of f to within $\frac{\varepsilon}{3t}$, provided we choose $\beta < 2^{-c \cdot \log(n/\varepsilon) \cdot \log \log n}$ for a large enough constant c . Consequently, by [Lemma 5.3.9](#), the composed restriction R preserves the expectation of f to within $\varepsilon/3$.

The next step is to show that read-once CNFs simplify under the composed restriction R . This is a straightforward consequence of our analysis in [Section 5.5.1](#). Indeed, for every k , each individual restriction $R^{(i)}$ is k -wise $(2\beta \cdot 2^{k/2})$ -close to $\mathcal{R}_{1/2}$ by [Theorem 2.3.2](#). Therefore, a composition of t restrictions such as $R^{(1)} \circ \dots \circ R^{(t)}$ is k -wise $(20\beta \cdot 2^{k/2})$ -close to \mathcal{R}_p where $p = 2^{-10}$. Therefore, by [Corollary 5.5.4](#), with probability $1 - \varepsilon/3$, the restricted function $f|_R$ is computable by a read-once CNF with at most m_* many clauses, where $m_* = \tilde{O}(\log w \cdot \log^2(1/\varepsilon))$, provided we choose $t = 10 \lceil \log w \rceil$ and $\beta < 2^{-c' \cdot (w + \log(1/\varepsilon))}$ for a large enough constant c' . Let E be the bad event that the restricted function is *not* computable by a read-once CNF with only m_* clauses.

The third step is to show that the small-bias distribution Z fools the simplified formula. This follows from our analysis in [Section 2.5.3](#). In that section, we argued that α -bias generators fool read-once CNFs with error $\exp(-\Omega(\log(1/\alpha)/\log n))$. Looking at the proof more closely, if the number of clauses is bounded by m_* , then the error is actually $\exp(-\Omega(\log(1/\alpha)/\log m_*))$. This error is at most $\varepsilon/3$ provided we choose a suitable value $\alpha = 2^{-\Theta(\log(1/\varepsilon) \cdot \log \log(w/\varepsilon))}$.

To wrap up the proof, let us compute the overall error of our PRG. Let U be a uniform random n -bit string. Then

$$\begin{aligned} |\mathbb{E}[f(R \circ Z)] - \mathbb{E}[f]| &\leq |\mathbb{E}[f(R \circ Z)] - \mathbb{E}[f(R \circ U)]| + |\mathbb{E}[f(R \circ U)] - \mathbb{E}[f]| \\ &\leq \mathbb{E}_R \left[\left| \mathbb{E}_Z[f|_R(Z)] - \mathbb{E}_U[f|_R(U)] \right| \right] + \varepsilon/3 \\ &\leq \mathbb{E}_R \left[\left| \mathbb{E}_Z[f|_R(Z)] - \mathbb{E}_U[f|_R(U)] \right| \mid \neg E \right] + 2\varepsilon/3 \\ &\leq \varepsilon. \end{aligned} \quad \square$$

To complete the proof of [Theorem 5.5.5](#), we need to eliminate the dependence on width from [Lemma 5.5.7](#). We accomplish this by a sandwiching argument.

Lemma 5.5.8. *Let f be a read-once CNF. For every $\varepsilon > 0$, f can be ε -sandwiched by read-once CNFs of width $\lceil \log(n/\varepsilon) \rceil$.*

Proof. Define f_u by deleting from f all clauses of width greater than $\lceil \log(n/\varepsilon) \rceil$, and meanwhile define f_ℓ by deleting all but the first $\lceil \log(n/\varepsilon) \rceil$ literals from each clause. Clearly, $f_\ell \leq f \leq f_u$. Furthermore, a clause of width $\lceil \log(n/\varepsilon) \rceil$ has expectation at least $1 - \varepsilon/n$. A read-once CNF can have at most n clauses, so by the union bound, $\Pr_x[f_u(x) \neq f_\ell(x)] \leq \varepsilon$. \square

[Theorem 5.5.5](#) is an immediate consequence of [Lemmas 2.5.2](#), [5.5.7](#) and [5.5.8](#).

5.5.3 Discussion: Two types of simplification

Looking again at the construction and analysis, our near-optimal PRG for read-once CNFs is based on combining *two distinct* simplification-under-restrictions lemmas:

- The first step of the PRG is applying a Forbes-Kelley pseudorandom restriction R (with \star -probability $p = 1/2$). As discussed in [Section 5.4.1](#), the fact that R preserves the expectation of every read-once CNF is equivalent to saying that every read-once CNF “simplifies on average” under a related partially-pseudorandom restriction (with \star -probability $1 - p = 1/2$).

- The second step of the PRG is arguing that read-once CNFs simplify with high probability under R^{ot} for some $t = O(\log \log n)$. Note that R^{ot} is a fully-pseudorandom restriction with \star -probability $1/\text{polylog}(n)$.

This is an attribute of the early termination framework more generally. In general, when we’re trying to prove the first simplification-under-restrictions lemma, two things are working in our favor: the relevant restriction is only *partially* pseudorandom, and it suffices to show simplification *on average*. On the other hand, when we’re trying to prove the second simplification-under-restrictions lemma, we have something else going for us: the \star -probability is relatively low. The early termination framework’s power comes from the fact that we get to *combine the advantages* of these two different settings.

In the decade since Gopalan, Meka, Reingold, Trevisan, and Vadhan introduced the early termination framework [GMRTV12], it has proven to be a versatile and powerful approach to PRG design, especially in the regime of near-optimal seed length [GMRTV12; MRT19; DHH19; Lee19; LV20; DHH20; DMRTV21]. Recently, Lyu introduced a different “partition-based” refinement of the iterated restrictions framework, which is also based on showing simplification under purely-pseudorandom restrictions, and used it to design an improved PRG for AC^0 circuits [Lyu22].

5.6 Fooling general branching programs via the IMZ framework

Let \mathcal{F} be a class of functions that we wish to fool. Let us suppose yet again that functions in \mathcal{F} simplify with high probability under restrictions. That is, for some values $p, \delta > 0$, we have

$$\Pr[f|_{\mathcal{R}_p} \in \mathcal{F}_{\text{simp}}] \geq 1 - \delta, \tag{5.10}$$

where $\mathcal{F}_{\text{simp}}$ is some class of “simpler” functions.

In Section 5.1, we used the polarizing walks framework to design a PRG for \mathcal{F} with a seed length of roughly $p^{-2} \cdot s$ (ignoring log factors), where s is the seed length of a PRG for $\mathcal{F}_{\text{simp}}$. Then, in Section 5.3, we showed that if it is possible to *partially derandomize* Equation (5.10), then we can improve the seed length to roughly $p^{-1} \cdot s$ using the Ajtai-Wigderson framework. In this section, we will present a framework due to Impagliazzo, Meka, and Zuckerman [IMZ19] (the “IMZ framework”). Assuming that it is possible to *fully derandomize* Equation (5.10), the IMZ framework gives a PRG for \mathcal{F} with a seed length of roughly $p^{-1} + r$, where r is the *description length* of functions in the “simpler” class $\mathcal{F}_{\text{simp}}$, i.e., $r = \log |\mathcal{F}_{\text{simp}}|$.

Observe that our measure of “simplicity” has changed. In all the previous sections, what mattered was the cost of *fooling* functions in $\mathcal{F}_{\text{simp}}$ (i.e., s), but now, what matters is the cost of *describing* functions in $\mathcal{F}_{\text{simp}}$ (i.e., r). Usually, the fooling cost s is much smaller than the description length r . (Indeed, ideally we hope for $s \approx \log r$; see Proposition 1.3.1.) Nevertheless, the IMZ framework is sometimes superior to the Ajtai-Wigderson framework, because the final seed length in the Ajtai-Wigderson framework is approximately the *product* $p^{-1} \cdot s$, whereas the final seed length in the IMZ framework is closer to the *sum* $p^{-1} + r$.

We emphasize that the IMZ framework requires a *fully-derandomized* simplification-under-restrictions lemma (just like the early termination framework that we discussed in Section 5.5). In Section 5.6.1, we will prove a fully-derandomized simplification-under-restrictions lemma for *general branching programs* where we only have a bound on the size of the program (i.e., the number of vertices). Then, in Section 5.6.2, we explain how to use such a lemma to construct a PRG.

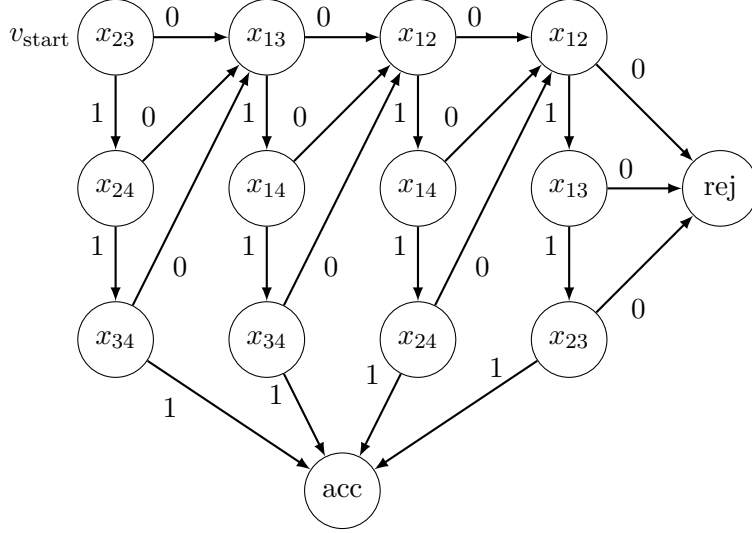


Figure 5.5: Let $n = m^2$, and let $x \in \{0, 1\}^n = \{0, 1\}^{m \times m}$ be the adjacency matrix of an undirected graph G . There is a branching program of size $O(m^3) = O(n^{1.5})$ that tests whether G has a triangle given x (the case $m = 4$ is shown above). In contrast, every *read-once* branching program computing this function must have width $2^{\Omega(n)}$, even if we allow arbitrary variable ordering. This follows from communication complexity lower bounds in the best-case partition model [PS84].

5.6.1 Shrinkage of branching programs under fully-pseudorandom restrictions

In this section, we study general size- m branching programs, with no restriction on the width or the number of times each variable is read.

Definition 5.6.1 (Unrestricted branching programs). *A size- m branching program over n input variables is a directed acyclic graph with at most m vertices. Each non-sink vertex v is labeled with an index $j_v \in [n]$ has two outgoing edges labeled 0 and 1. A subset V_{accept} of the sink vertices are designated as “accepting vertices.” Given an input $x \in \{0, 1\}^n$, the program starts at a designated “start vertex” v_{start} , and in each step, having reached a vertex v , the program queries x_{j_v} and traverses the corresponding outgoing edge. Eventually, the program reaches a sink vertex v , and $f(x) = 1 \iff v \in V_{\text{accept}}$. (See Figure 5.5.)*

Let $\text{BP}(f)$ denote the size of the smallest branching program computing f . One can easily show that branching programs *shrink* under truly random restrictions, in the sense that $\mathbb{E}[\text{BP}(f|_{\mathcal{R}_p})] \leq p \cdot \text{BP}(f)$. We will prove that similar shrinkage occurs *with high probability* rather than in expectation, and furthermore that it occurs under a *fully pseudorandom restriction*. In particular, our pseudorandom restriction distribution is as follows.

Definition 5.6.2 (k -wise independent restrictions). *Let R be a distribution over $\{0, 1, \star\}^n$. We say that R is a k -wise independent p -regular restriction if the coordinates of R are k -wise independent, and the marginal distributions are given by*

$$R_i = \begin{cases} \star & \text{with probability } p \\ 0 & \text{with probability } (1-p)/2 \\ 1 & \text{with probability } (1-p)/2. \end{cases}$$

Equivalently, R is k -wise 0-close to R_p (see Definition 5.5.1).

Lemma 5.6.3 (High-probability shrinkage of branching programs under fully-pseudorandom restrictions). *For every $\delta > 0$, there is a value $k = O(\log(1/\delta))$ such that for every $f: \{0, 1\}^n \rightarrow \{0, 1\}$ and every $p > 0$, if $R \in \{0, 1, \star\}^n$ is a k -wise independent p -regular restriction, then*

$$\Pr \left[\text{BP}(f|_R) \leq \lceil p \cdot \text{BP}(f) \rceil \cdot 2^{O(\sqrt{\log(1/\delta)})} \right] \geq 1 - \delta.$$

As discussed previously, what we really care about is the *description length* of $f|_R$, but this can be bounded in terms of the branching program size $\text{BP}(f|_R)$. To prove [Lemma 5.6.3](#), we rely on a tail bound for sums of k -wise independent random variables, which we cite without proof.¹⁰

Theorem 5.6.4 (Tail bound for sums of k -wise independent random variables [[BR94](#), Lemma 2.3]). *Let $X_1, \dots, X_n \in [0, 1]$ be k -wise independent, where $k \geq 4$ is an even integer, and let $X = \sum_{i=1}^n X_i$. Then for any $\Delta > 0$,*

$$\Pr [|X - \mathbb{E}[X]| \geq \Delta] \leq 8 \cdot \left(\frac{k \mathbb{E}[X] + k^2}{\Delta^2} \right)^{k/2}.$$

Proof of [Lemma 5.6.3](#). Identify f with a branching program computing f of size $\text{BP}(f)$. For $i \in [n]$, let m_i be the number of nodes in f that query x_i . Let H be the set of “heavy variables,” namely $H = \{i : m_i > h\}$ where $h = p \cdot \text{BP}(f) \cdot 2\sqrt{\log(1/\delta)}$. We first show that with high probability, few heavy variables are left alive. Indeed, since $k > \sqrt{\log(1/\delta)}$, we have

$$\begin{aligned} \Pr \left[|H \cap R^{-1}(\star)| \geq \sqrt{\log(1/\delta)} \right] &\leq \left(\frac{|H|}{\sqrt{\log(1/\delta)}} \right) \cdot p^{\sqrt{\log(1/\delta)}} \leq (p|H|)^{\sqrt{\log(1/\delta)}} \\ &\leq (p \cdot \text{BP}(f)/h)^{\sqrt{\log(1/\delta)}} \\ &= \delta. \end{aligned}$$

Now let us consider the “light” variables. For each $i \notin H$, let $X_i = m_i \cdot \mathbb{1}[R_i = \star]/h \in [0, 1]$. Let $X = \sum_{i \in H} X_i$, so $\mathbb{E}[X] \leq p \cdot \text{BP}(f)/h < 1$. By [Theorem 5.6.4](#) with $\Delta = O(\log(1/\delta))$, we have

$$\Pr[X \leq O(\log(1/\delta))] \geq 1 - \delta.$$

Our branching program for $f|_R$ begins by querying all the variables in $H \cap R^{-1}(\star)$ and storing all those values in memory. Then it simulates the branching program for f , skipping queries to variables in $H \cup R^{-1}(\{0, 1\})$ since those values are known. The size of the branching program is

$$2^{|H \cap R^{-1}(\star)|} - 1 + 2^{|H \cap R^{-1}(\star)|} \cdot \sum_{i \in R^{-1}(\star) \setminus H} m_i < (1 + h \cdot X) \cdot 2^{|H \cap R^{-1}(\star)|},$$

which is bounded by $\lceil p \cdot \text{BP}(f) \rceil \cdot 2^{O(\sqrt{\log(1/\delta)})}$ except with probability 2δ . Replacing δ with $\delta/2$ completes the proof. \square

The parameters of [Lemma 5.6.3](#) are perhaps a bit disappointing. Let $m = \text{BP}(f)$. When $\delta < 2^{-\Theta(\log^2 m)}$, the lemma breaks down: it is not able to show that *any* shrinkage occurs with probability $1 - \delta$. Unfortunately, this is unavoidable. Indeed, by a standard counting argument,

¹⁰Here we are citing a bound due to Bellare and Rompel [[BR94](#)]. Skorski has shown an improvement to Bellare and Rompel’s bound [[Sko22](#)], but the improved bound is slightly more complicated and it makes no difference in our application, so we stick with Bellare and Rompel’s simpler bound [[BR94](#)].

there exists a function f with $\text{BP}(f) \geq m$ that only reads $k \stackrel{\text{def}}{=} O(\log m)$ of the input variables. For this function f , assuming $p \geq 1/m$, even under a truly random restriction $R = R_p$, we have

$$\Pr[\text{BP}(f|_R) = \text{BP}(f)] \geq p^k \geq 2^{-O(\log^2 m)}.$$

Thus, in the regime $p \geq 1/m$ (which is the most interesting regime), one cannot prove a shrinkage lemma where the failure probability is exponentially small compared to m .

5.6.2 PRGs from fully-derandomized shrinkage lemmas

Now let us present the IMZ reduction.

Theorem 5.6.5 (Simplification under fully-pseudorandom restrictions \implies PRG [IMZ19]). *Let \mathcal{F} and $\mathcal{F}_{\text{simp}}$ be classes of functions $f: \{0, 1\}^n \rightarrow \{0, 1\}$. Assume that \mathcal{F} is closed under restrictions and shifts, and assume that $\mathcal{F}_{\text{simp}}$ contains the constant 0 function. Let $\delta > 0$, and let R be a random variable over $\{0, 1, \star\}^n$ that can be explicitly sampled using q truly random bits such that*

$$\forall f \in \mathcal{F}, \Pr[f|_R \in \mathcal{F}_{\text{simp}}] \geq 1 - \delta.$$

Assume that we can explicitly compute a value p such that for every $i \in [n]$, we have $\Pr[R_i = \star] \geq p$. Let $r \in \mathbb{N}$, and assume that (a) $\log |\mathcal{F}_{\text{simp}}| \leq r$ and (b) there is an explicit δ -PRG for $\mathcal{F}_{\text{simp}}$ with seed length r .¹¹ Then there is an explicit PRG that fools \mathcal{F} with error $O(t\delta)$ and seed length $O(t \cdot (q + \log(r/\delta)) + r)$, where $t = \lceil p^{-1} \ln(n/\delta) \rceil$.

We think of q and $\text{polylog}(nm/\delta)$ as “small,” so the seed length in [Theorem 5.6.5](#) is indeed approximately $p^{-1} + r$, as suggested at the beginning of this section. For branching programs, [Theorem 5.6.5](#) implies the following PRG.

Corollary 5.6.6 (PRG for branching programs [IMZ19]). *For any n, m, ε , there is an explicit ε -PRG for size- m branching programs with seed length $\sqrt{m} \cdot 2^{O(\sqrt{\log(m/\varepsilon)})} \cdot \text{polylog } n$.*

When $\varepsilon = 1/\text{poly}(m)$ and $m \geq n$, the seed length in [Corollary 5.6.6](#) is $m^{\frac{1}{2}+o(1)}$.

Proof sketch. Assume without loss of generality that $\log n \leq m \leq n^2$. Let p be the largest power of two with $p < 1/\sqrt{m}$. Let $t = \lceil p^{-1} \ln(n/\varepsilon) \rceil$ and let $\delta = \Theta(\varepsilon/t)$. Let R be a k -wise independent p -regular restriction where $k = O(\log(1/\delta))$. By a construction similar to the proof of [Theorem 2.1.3](#), the distribution R can be sampled explicitly using q truly random bits, where

$$q = O(k \log(n/p)) = O(\log(m/\varepsilon) \log n).$$

Let $\mathcal{F}_{\text{simp}}$ be the class of all branching programs of size at most m' , where

$$m' = \lceil p \cdot m \rceil \cdot 2^{O(\sqrt{\log(1/\delta)})} = \sqrt{m} \cdot 2^{O(\sqrt{\log(m/\varepsilon)})}.$$

By [Lemma 5.6.3](#), for every size- m branching program f , we have $\Pr[f|_R \in \mathcal{F}_{\text{simp}}] \geq 1 - \delta$. Let $r = O(m' \log n)$. Then $\log |\mathcal{F}_{\text{simp}}| \leq r$, and furthermore there is an explicit PRG that perfectly fools $\mathcal{F}_{\text{simp}}$ with seed length r , namely an m' -wise uniform generator. Therefore, by [Theorem 5.6.5](#), there is an explicit PRG for size- m branching programs with error $O(t\delta) = \varepsilon$ and seed length

$$O(t \cdot (s + \log(r/\delta)) + r) = \sqrt{m} \cdot 2^{O(\sqrt{\log(m/\varepsilon)})} \cdot \text{polylog}(n). \quad \square$$

¹¹Given condition (a), condition (b) is relatively mild; note that the optimal seed length would be $O(\log(r/\delta))$ (see [Proposition 1.3.1](#)).

In their original paper, Impagliazzo, Meka, and Zuckerman used the IMZ framework to design PRGs for a few additional classes, such as De Morgan formulas [IMZ19]. Later, Hatami, Hoza, Tal, and Tell gave improved PRGs for branching programs and De Morgan formulas using variants of the IMZ framework [HHTT22]. For branching programs, the improved seed length is $\sqrt{m} \cdot \text{polylog}(n/\varepsilon)$, which is close to the lack-of-lower-bounds barrier. See also the work of Cheraghchi, Kabanets, Lu, and Myrasiotis for another variation on the IMZ framework [CKLM20].

Now let us get started proving the basic IMZ reduction (Theorem 5.6.5). The proof draws inspiration from the Nisan-Zuckerman generator (Section 3.4). The high-level intuition is that when we do a restriction, the restricted function cannot encode much information about the random bits we have used so far (since it can be succinctly described), and therefore we can use an extractor to recycle the random bits.

In detail, let $G: \{0, 1\}^r \rightarrow \{0, 1\}^n$ be a δ -PRG for $\mathcal{F}_{\text{simp}}$, let $\text{Ext}: \{0, 1\}^\ell \times \{0, 1\}^d \rightarrow \{0, 1\}^r$ be an $(\ell - r, \delta)$ -extractor, and let G' denote the following PRG.

1. Sample t independent copies $R^{(1)}, \dots, R^{(t)}$ of the restriction R .
2. Sample $X, Y^{(1)}, \dots, Y^{(t)}$ uniformly at random, let $Z^{(i)} = \text{Ext}(X, Y^{(i)})$, and output

$$\sum_{i=1}^t R^{(i)} \circ G(Z^{(i)}),$$

where the sum is over \mathbb{F}_2^n .

Lemma 5.6.7 (Correctness of the IMZ reduction). *Under the assumptions of Theorem 5.6.5, the generator G' defined above fools \mathcal{F} with error $O(t\delta)$.*

Proof. The proof is a hybrid argument. Sample $U^{(1)}, \dots, U^{(t)} \in \{0, 1\}^n$ independently and uniformly at random. Define hybrid distributions $H^{(0)}, \dots, H^{(t)}$ by

$$H^{(j)} = \left(\sum_{i=1}^j R^{(i)} \circ U^{(i)} \right) + \sum_{i=j+1}^t R^{(i)} \circ G(Z^{(i)}),$$

i.e., in the first j terms of the sum, we fill in the stars of $R^{(i)}$ using truly random bits instead of the pseudorandom bits $G(Z^{(i)})$. Fix some $f \in \mathcal{F}$ and some $j \in [t]$. Let us show that $\mathbb{E}[f(H^{(j-1)})] \approx \mathbb{E}[f(H^{(j)})]$.

There is a random variable B , independent of $U^{(j)}$ and $Y^{(j)}$, such that

$$\begin{aligned} H^{(j-1)} &= B + R^{(j)} \circ G(Z^{(j)}) \\ H^{(j)} &= B + R^{(j)} \circ U^{(j)}. \end{aligned}$$

(Note that B and X are *not* independent.) Define $f^{+B}(x) = f(x + B)$, define $F = f^{+B}|_{R^{(j)}}$, and define

$$F' = \begin{cases} F & \text{if } F \in \mathcal{F}_{\text{simp}} \\ \text{the 0 function} & \text{if } F \notin \mathcal{F}_{\text{simp}}. \end{cases}$$

By construction, $F' \in \mathcal{F}_{\text{simp}}$, so F' can be described using r bits. By Lemma 3.4.6, $\tilde{H}_{\min}(X | F') \geq \ell - r$. Therefore, by Lemma 3.4.7,

$$d_{\text{TV}}((Z^{(j)}, F'), (U, F')) \leq 3\delta,$$

where U is a uniform random r -bit string independent of F' . Applying a deterministic function can only make two distributions closer, so

$$|\mathbb{E}[F'(G(Z^{(j)}))] - \mathbb{E}[F'(G(U))]| \leq 3\delta.$$

Since $F' \in \mathcal{F}_{\text{simp}}$, the generator G fools F' with error δ , hence $|\mathbb{E}[F'(G(U))] - \mathbb{E}[F'(U^{(j)})]| \leq \delta$. Furthermore, with probability $1 - \delta$, we have $F \equiv F'$. Therefore, overall,

$$|\mathbb{E}[F(G(Z^{(j)}))] - \mathbb{E}[F(U^{(j)})]| \leq 5\delta,$$

or equivalently,

$$|\mathbb{E}[f(H^{(j-1)})] - \mathbb{E}[f(H^{(j)})]| \leq 5\delta.$$

Clearly, $H^{(0)}$ is the output distribution of our PRG G' . By the triangle inequality, $H^{(0)}$ and $H^{(t)}$ are nearly indistinguishable to f . To complete the proof, let us show that $H^{(t)}$ is statistically close to uniform. Indeed, for each $j \in [n]$, we have

$$\Pr[\forall i \in [t], R_j^{(i)} \neq \star] \leq (1 - p)^t \leq e^{-pt} \leq \delta/n.$$

Therefore, by the union bound, with probability at least $1 - \delta$, we have $\bigcup_{i=1}^t (R^{(i)})^{-1}(\star) = [n]$. In this case, with respect to the randomness of $U^{(1)}, \dots, U^{(t)}$, the distribution $H^{(t)}$ is uniform. Therefore, overall, the total variation distance between $H^{(t)}$ and the uniform distribution is at most δ , and hence G' fools f with error $(5t + 1)\delta$. \square

To complete the proof of [Theorem 5.6.5](#), we bound the seed length of G' .

Proof of [Theorem 5.6.5](#). The restrictions $R^{(1)}, \dots, R^{(t)}$ cost qt truly random bits in total. Using the GUV extractor ([Theorem 3.4.4](#)), the source length ℓ of the extractor Ext is $O(r)$, and its seed length is $d = O(\log(r/\delta))$. Therefore, the total seed length is $(q + d)t + \ell$, which is $O((q + \log(r/\delta)) \cdot t + r)$ as claimed. \square

Chapter 6

Table of PRGs

For reference, we conclude this text with a table of the best explicit PRG constructions currently known, arranged by the model they fool. The table is not meant to be exhaustive; only a selection of important computational models are included. In each case, we only record a single state-of-the-art seed length, which in many cases is superior to the PRG constructions that we presented.

6.1 Circuit models

In the table below, we use d to denote depth and m to denote size. Assume $d = O(1)$ and $m \geq n$.

Model	Seed length	Approach	Reference
Depth-1 \mathbf{AC}^0 circuits	$O(\log(1/\varepsilon) + \log \log n)$	k -wise δ -bias	Folklore
\mathbf{AC}^0 circuits	$\tilde{O}(\log^{d-1} m \cdot \log(m/\varepsilon))$	Variant of the Ajtai-Wigderson framework	[Lyu22]
Read-once CNFs/DNFs	$O(\log n) + \tilde{O}(\log(1/\varepsilon))$	Iterated restrictions with early termination	[DHH20]
Read-once \mathbf{AC}^0	$\tilde{O}(\log(n/\varepsilon))$	Iterated restrictions	[DHH19; DMRTV21]
De Morgan formulas	$m^{1/3+o(1)} \cdot \text{polylog}(1/\varepsilon)$	Variant of the IMZ framework	[HHTT22]
Read-once De Morgan formulas	$O(\log^2 n \cdot \log(n/\varepsilon))$	Iterated restrictions	[FK18]

6.2 Branching program models

In the table below, we use m to denote size and w to denote width. Assume $m \geq n$.

Model	Seed length	Approach	Reference
Unrestricted branching programs	$\sqrt{m} \cdot \text{polylog}(n/\varepsilon)$	Variant of the IMZ framework	[HHTT22]
Width-2 branching programs that read d bits at a time	$O(d \log n + d \cdot 2^d \cdot \log(m/\varepsilon))$	Sum of d δ -biased distributions	[BDVY13]
Standard-order ROBPs with $w = 3$	$\tilde{O}(\log n \cdot \log(1/\varepsilon))$	Iterated restrictions with early termination	[MRT19]
Standard-order ROBPs with $4 \leq w \leq n$	$O(\log(n/\varepsilon) \cdot \log n)$	Recycling seeds	[Nis92; INW94]
Standard-order ROBPs with $w \gg n$	$O\left(\frac{\log(w/\varepsilon) \cdot \log n}{\log \log w}\right)$	Recycling seeds	[Arm98; KNW08]
Standard-order regular ROBPs	$\tilde{O}(\log(w/\varepsilon) \cdot \log n)$	INW generator	[BRRY14]
Standard-order permutation ROBPs with $w = O(1)$	$O(\log n \cdot \log(1/\varepsilon))$	INW generator	[KNP11; De11; Ste12]
Arbitrary-order ROBPs	$O(\log(wn/\varepsilon) \cdot \log^2 n)$	Iterated restrictions	[FK18]
Arbitrary-order ROBPs with $w = O(1)$	$\tilde{O}(\log(n/\varepsilon) \cdot \log n)$	Iterated restrictions	[FK18]
Arbitrary-order permutation ROBPs with $w = O(1)$	$\tilde{O}(\log n \cdot \log(1/\varepsilon))$	Polarizing random walks	[CHHL19]
Decision trees, or more generally parity decision trees	$O(\log(m/\varepsilon))$	δ -bias	[KM93]

6.3 Algebraic models

Model	Seed length	Approach	Reference
Parity functions	$O(\log(n/\varepsilon))$	Balanced codes	[NN93]
Parities of at most k bits	$O(\log(k/\varepsilon) + \log \log n)$	ε -biased seed for k -wise uniform generator	[NN93]
Degree- d polynomials over \mathbb{F}_2	$O(d \log n + d2^d \log(1/\varepsilon))$	Sum of d δ -biased distributions	[Vio09]

6.4 Models based on locality

Model	Seed length	Approach	Reference
$[-1, 1]$ -valued k -juntas	$O(k + \log(1/\varepsilon) + \log \log n)$	k -wise δ -bias	[NN93]
Two-dimensional combinatorial rectangles	$\frac{n}{2} + O(\log(1/\varepsilon))$	Random edge of expander	[INW94]
d -dimensional combinatorial rectangles	$\tilde{O}(n/d + \log(1/\varepsilon) + \log \log n)$	Iterative alphabet reduction	[GY20]
Two-party communication protocols with cost m	$\frac{n}{2} + O(m + \log(1/\varepsilon))$	Random edge of expander	[INW94]

Bibliography

- [Aar10] Scott Aaronson. “BQP and the polynomial hierarchy”. In: *Proc. 42nd Annual ACM Symposium on Theory of Computing (STOC)*. 2010, pp. 141–150. DOI: [10.1145/1806689.1806711](https://doi.org/10.1145/1806689.1806711) (cit. on p. 21).
- [AB09] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009. DOI: [10.1017/CB09780511804090](https://doi.org/10.1017/CB09780511804090) (cit. on p. 9).
- [ABI86] Noga Alon, László Babai, and Alon Itai. “A fast and simple randomized parallel algorithm for the maximal independent set problem”. In: *J. Algorithms* 7.4 (1986), pp. 567–583. ISSN: 0196-6774. DOI: [10.1016/0196-6774\(86\)90019-2](https://doi.org/10.1016/0196-6774(86)90019-2) (cit. on pp. 15, 16).
- [ACGS88] Werner Alexi, Benny Chor, Oded Goldreich, and Claus P. Schnorr. “RSA and Rabin functions: certain parts are as hard as the whole”. In: *SIAM J. Comput.* 17.2 (1988). Special issue on cryptography, pp. 194–209. ISSN: 0097-5397. DOI: [10.1137/0217013](https://doi.org/10.1137/0217013) (cit. on p. 69).
- [ACR98] Alexander E. Andreev, Andrea E. F. Clementi, and José D. P. Rolim. “A new general derandomization method”. In: *J. ACM* 45.1 (1998), pp. 179–213. ISSN: 0004-5411. DOI: [10.1145/273865.273933](https://doi.org/10.1145/273865.273933) (cit. on p. 12).
- [ACRT99] Alexander E. Andreev, Andrea E. F. Clementi, José D. P. Rolim, and Luca Trevisan. “Weak random sources, hitting sets, and BPP simulations”. In: *SIAM J. Comput.* 28.6 (1999), pp. 2103–2116. ISSN: 0097-5397. DOI: [10.1137/S0097539797325636](https://doi.org/10.1137/S0097539797325636) (cit. on p. 12).
- [AGHP92] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. “Simple Constructions of Almost k -wise Independent Random Variables”. In: *Random Structures & Algorithms* 3.3 (1992), pp. 289–304. DOI: [10.1002/rsa.3240030308](https://doi.org/10.1002/rsa.3240030308) (cit. on pp. 19, 20).
- [AKS87] Miklós Ajtai, János Komlós, and Endre Szemerédi. “Deterministic Simulation in LOGSPACE”. In: *Proc. 19th Annual ACM Symposium on Theory of Computing (STOC)*. 1987, pp. 132–140. DOI: [10.1145/28395.28410](https://doi.org/10.1145/28395.28410) (cit. on pp. 12, 60).
- [Alo21] Noga Alon. “Explicit expanders of every degree and size”. In: *Combinatorica* 41.4 (2021), pp. 447–463. ISSN: 0209-9683. DOI: [10.1007/s00493-020-4429-x](https://doi.org/10.1007/s00493-020-4429-x) (cit. on p. 46).
- [Alo86] Noga Alon. “Eigenvalues and expanders”. In: *Combinatorica* 6.2 (1986), pp. 83–96. ISSN: 0209-9683. DOI: [10.1007/BF02579166](https://doi.org/10.1007/BF02579166) (cit. on p. 46).

- [Arm98] Roy Armoni. “On the derandomization of space-bounded computations”. In: *Proc. 2nd International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 1998, pp. 47–59. DOI: [10.1007/3-540-49543-6_5](https://doi.org/10.1007/3-540-49543-6_5) (cit. on pp. 59, 109).
- [ASWZ96] Roy Armoni, Michael Saks, Avi Wigderson, and Shiyu Zhou. “Discrepancy sets and pseudorandom generators for combinatorial rectangles”. In: *Proc. 37th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 1996, pp. 412–421. DOI: [10.1109/SFCS.1996.548500](https://doi.org/10.1109/SFCS.1996.548500) (cit. on p. 47).
- [AW89] Miklós Ajtai and Avi Wigderson. “Deterministic Simulation of Probabilistic Constant-Depth Circuits”. In: *Advances in Computing Research – Randomness and Computation* 5 (1989), pp. 199–23 (cit. on pp. 88, 89).
- [Baz09] Louay M. J. Bazzi. “Polylogarithmic Independence Can Fool DNF Formulas”. In: *SIAM J. Comput.* 38.6 (2009), pp. 2220–2272. DOI: [10.1137/070691954](https://doi.org/10.1137/070691954) (cit. on pp. 35, 37, 39).
- [BBS86] L. Blum, M. Blum, and M. Shub. “A simple unpredictable pseudorandom number generator”. In: *SIAM J. Comput.* 15.2 (1986), pp. 364–383. ISSN: 0097-5397. DOI: [10.1137/0215025](https://doi.org/10.1137/0215025) (cit. on pp. 7, 69).
- [BCG20] Mark Braverman, Gil Cohen, and Sumegha Garg. “Pseudorandom pseudo-distributions with near-optimal error for read-once branching programs”. In: *SIAM J. Comput.* 49.5 (2020), STOC18–242–STOC18–299. ISSN: 0097-5397. DOI: [10.1137/18M1197734](https://doi.org/10.1137/18M1197734) (cit. on pp. 12, 51).
- [BDFP86] Allan Borodin, Danny Dolev, Faith E. Fich, and Wolfgang Paul. “Bounds for width two branching programs”. In: *SIAM J. Comput.* 15.2 (1986), pp. 549–560. ISSN: 0097-5397. DOI: [10.1137/0215040](https://doi.org/10.1137/0215040) (cit. on pp. 24, 25).
- [BDVY13] Andrej Bogdanov, Zeev Dvir, Elad Verbin, and Amir Yehudayoff. “Pseudorandomness for Width-2 Branching Programs”. In: *Theory of Computing* 9 (2013), pp. 283–293. DOI: [10.4086/toc.2013.v009a007](https://doi.org/10.4086/toc.2013.v009a007) (cit. on pp. 24, 25, 30, 109).
- [BF99] Harry Buhrman and Lance Fortnow. “One-Sided Versus Two-Sided Error in Probabilistic Computation”. In: *Proc. 16th Symposium on Theoretical Aspects of Computer Science (STACS)*. 1999, pp. 100–109. ISBN: 978-3-540-49116-3. DOI: [10.1007/3-540-49116-3_9](https://doi.org/10.1007/3-540-49116-3_9) (cit. on p. 12).
- [BFNW93] László Babai, Lance Fortnow, Noam Nisan, and Avi Wigderson. “BPP has subexponential time simulations unless EXPTIME has publishable proofs”. In: *Comput. Complexity* 3.4 (1993), pp. 307–318. ISSN: 1016-3328. DOI: [10.1007/BF01275486](https://doi.org/10.1007/BF01275486) (cit. on p. 69).
- [BHPP22] Andrej Bogdanov, William M. Hoza, Gautam Prakriya, and Edward Pyne. “Hitting Sets for Regular Branching Programs”. In: *Proc. 37th Computational Complexity Conference (CCC)*. 2022, 3:1–3:22. ISBN: 978-3-95977-241-9. DOI: [10.4230/LIPIcs.CCC.2022.3](https://doi.org/10.4230/LIPIcs.CCC.2022.3) (cit. on pp. 51, 55).
- [BHST87] László Babai, Péter Hajnal, Endre Szemerédi, and György Turán. “A lower bound for read-once-only branching programs”. In: *J. Comput. System Sci.* 35.2 (1987), pp. 153–162. ISSN: 0022-0000. DOI: [10.1016/0022-0000\(87\)90010-9](https://doi.org/10.1016/0022-0000(87)90010-9) (cit. on p. 49).

- [BJS01] Paul Beame, T. S. Jayram, and Michael Saks. “Time-space tradeoffs for branching programs”. In: *J. Comput. System Sci.* 63.4 (2001), pp. 542–572. ISSN: 0022-0000. DOI: [10.1006/jcss.2001.1778](https://doi.org/10.1006/jcss.2001.1778) (cit. on p. 49).
- [BLP15] Paul Beame, Vincent Liew, and Mihai Pătraşcu. “Finding the median (obliviously) with bounded space”. In: *Proc. 42nd International Colloquium on Automata, Languages and Programming (ICALP)*. 2015, pp. 103–115. DOI: [10.1007/978-3-662-47672-7_9](https://doi.org/10.1007/978-3-662-47672-7_9) (cit. on p. 49).
- [BM84] Manuel Blum and Silvio Micali. “How to Generate Cryptographically Strong Sequences of Pseudorandom Bits”. In: *SIAM J. Comput.* 13.4 (1984), pp. 850–864. ISSN: 0097-5397. DOI: [10.1137/0213053](https://doi.org/10.1137/0213053) (cit. on pp. 7, 65, 69).
- [BNS92] László Babai, Noam Nisan, and Máriaó Szegedy. “Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs”. In: *J. Comput. System Sci.* 45.2 (1992), pp. 204–232. ISSN: 0022-0000. DOI: [10.1016/0022-0000\(92\)90047-M](https://doi.org/10.1016/0022-0000(92)90047-M) (cit. on p. 68).
- [Bog05] Andrej Bogdanov. “Pseudorandom generators for low degree polynomials”. In: *Proc. 37th Annual ACM Symposium on Theory of Computing (STOC)*. 2005, pp. 21–30. DOI: [10.1145/1060590.1060594](https://doi.org/10.1145/1060590.1060594) (cit. on p. 26).
- [Bop97] Ravi B. Boppana. “The Average Sensitivity of Bounded-Depth Circuits”. In: *Inf. Process. Lett.* 63.5 (1997), pp. 257–261. DOI: [10.1016/S0020-0190\(97\)00131-2](https://doi.org/10.1016/S0020-0190(97)00131-2) (cit. on p. 38).
- [Bor20] Charles Bordenave. “A new proof of Friedman’s second eigenvalue theorem and its extension to random lifts”. In: *Ann. Sci. Éc. Norm. Supér. (4)* 53.6 (2020), pp. 1393–1439. ISSN: 0012-9593. DOI: [10.24033/asens.245](https://doi.org/10.24033/asens.245) (cit. on p. 46).
- [BPW11] Andrej Bogdanov, Periklis A. Papakonstantinou, and Andrew Wan. “Pseudorandomness for Read-Once Formulas”. In: *FOCS*. Ed. by Rafail Ostrovsky. IEEE, 2011, pp. 240–246. ISBN: 978-1-4577-1843-4 (cit. on pp. 85, 94).
- [BR94] M. Bellare and J. Rompel. “Randomness-efficient oblivious sampling”. In: *Proc. 35th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 1994, pp. 276–287. DOI: [10.1109/SFCS.1994.365687](https://doi.org/10.1109/SFCS.1994.365687) (cit. on p. 104).
- [Bra10] Mark Braverman. “Polylogarithmic independence fools AC⁰ circuits”. In: *J. ACM* 57.5 (2010), 28:1–28:10. ISSN: 0004-5411. DOI: [10.1145/1754399.1754401](https://doi.org/10.1145/1754399.1754401) (cit. on pp. 9, 37, 38, 41, 42).
- [BRRY14] Mark Braverman, Anup Rao, Ran Raz, and Amir Yehudayoff. “Pseudorandom generators for regular branching programs”. In: *SIAM J. Comput.* 43.3 (2014), pp. 973–986. ISSN: 0097-5397. DOI: [10.1137/120875673](https://doi.org/10.1137/120875673) (cit. on pp. 51–53, 55, 109).
- [BRS91] Richard Beigel, Nick Reingold, and Daniel A. Spielman. “The perceptron strikes back”. In: *Proc. 6th Annual IEEE Conference on Structure in Complexity Theory*. 1991, pp. 286,287,288,289,290,291. DOI: [10.1109/SCT.1991.160270](https://doi.org/10.1109/SCT.1991.160270) (cit. on pp. 41, 42).
- [BT11] Avraham Ben-Aroya and Amnon Ta-Shma. “A combinatorial construction of almost-Ramanujan graphs using the zig-zag product”. In: *SIAM J. Comput.* 40.2 (2011), pp. 267–290. ISSN: 0097-5397. DOI: [10.1137/080732651](https://doi.org/10.1137/080732651) (cit. on p. 46).

- [BV10] Joshua Brody and Elad Verbin. “The Coin Problem and Pseudorandomness for Branching Programs”. In: *Proc. 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2010, pp. 30–39. DOI: [10.1109/FOCS.2010.10](https://doi.org/10.1109/FOCS.2010.10). URL: <https://www.cs.swarthmore.edu/~brody/papers/TheCoinProblemFOCS.pdf> (cit. on pp. 55, 83).
- [CDRST21] Gil Cohen, Dean Doron, Oren Renard, Ori Sberlo, and Amnon Ta-Shma. “Error reduction for weighted PRGs against read once branching programs”. In: *Proc. 36th Computational Complexity Conference (CCC)*. 2021, 22:1–22:17. DOI: [10.4230/LIPIcs.CCC.2021.22](https://doi.org/10.4230/LIPIcs.CCC.2021.22) (cit. on pp. 12, 51).
- [CG88] Benny Chor and Oded Goldreich. “Unbiased Bits from Sources of Weak Randomness and Probabilistic Communication Complexity”. In: *SIAM J. on Computing* 17.2 (1988), pp. 230–261. DOI: [10.1137/0217015](https://doi.org/10.1137/0217015) (cit. on p. 56).
- [CGHFRS85] Benny Chor, Oded Goldreich, Johan Håstad, Joel Friedman, Steven Rudich, and Roman Smolensky. “The Bit Extraction Problem or t -Resilient Functions”. In: *Proc. 26th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 1985, pp. 396–407. DOI: [10.1109/SFCS.1985.55](https://doi.org/10.1109/SFCS.1985.55) (cit. on pp. 15, 16).
- [CGLLS21] Eshan Chattopadhyay, Jason Gaitonde, Chin Ho Lee, Shachar Lovett, and Abhishek Shetty. “Fractional Pseudorandom Generators from Any Fourier Level”. In: *Proc. 36th Computational Complexity Conference (CCC)*. 2021, 10:1–10:24. ISBN: 978-3-95977-193-1. DOI: [10.4230/LIPIcs.CCC.2021.10](https://doi.org/10.4230/LIPIcs.CCC.2021.10) (cit. on p. 88).
- [CH22] Kuan Cheng and William M. Hoza. “Hitting Sets Give Two-Sided Derandomization of Small Space”. In: *Theory of Computing* 18.21 (2022), pp. 1–32. DOI: [10.4086/toc.2022.v018a021](https://doi.org/10.4086/toc.2022.v018a021) (cit. on p. 12).
- [Cha18] Eshan Chattopadhyay. *Pseudorandomness and Combinatorial Constructions*. Course at Cornell University. 2018. URL: <https://courses.cs.cornell.edu/cs6815/2018fa/> (cit. on p. 9).
- [Cha19] Eshan Chattopadhyay. *Pseudorandomness and Combinatorial Constructions*. Course at Cornell University. 2019. URL: <https://courses.cs.cornell.edu/cs6815/2019fa/> (cit. on p. 9).
- [Cha22] Eshan Chattopadhyay. *Pseudorandomness and Combinatorial Constructions*. Course at Cornell University. 2022. URL: <https://courses.cs.cornell.edu/cs6815/2022fa/> (cit. on p. 9).
- [CHHL19] Eshan Chattopadhyay, Pooya Hatami, Kaave Hosseini, and Shachar Lovett. “Pseudorandom generators from polarizing random walks”. In: *Theory Comput.* 15.1 (2019), pp. 1–26. DOI: [10.4086/toc.2019.v015a010](https://doi.org/10.4086/toc.2019.v015a010) (cit. on pp. 9, 55, 72, 74, 76, 78, 86, 109).
- [CHLT18] Eshan Chattopadhyay, Pooya Hatami, Shachar Lovett, and Avishay Tal. “Pseudorandom Generators from the Second Fourier Level and Applications to AC0 with Parity Gates”. In: *Proc. 10th Conference on Innovations in Theoretical Computer Science (ITCS)*. 2018, 22:1–22:15. ISBN: 978-3-95977-095-8. DOI: [10.4230/LIPIcs.ITCS.2019.22](https://doi.org/10.4230/LIPIcs.ITCS.2019.22) (cit. on p. 88).
- [CHRT18] Eshan Chattopadhyay, Pooya Hatami, Omer Reingold, and Avishay Tal. “Improved pseudorandomness for unordered branching programs through local monotonicity”. In: *Proc. 50th Annual ACM Symposium on Theory of Computing (STOC)*. 2018, pp. 363–375. DOI: [10.1145/3188745.3188800](https://doi.org/10.1145/3188745.3188800) (cit. on pp. 81, 83, 87).

- [CIS18] Marco L. Carmosino, Russell Impagliazzo, and Manuel Sabin. “Fine-Grained Derandomization: From Problem-Centric to Resource-Centric Complexity”. In: *Proc. 45th International Colloquium on Automata, Languages and Programming (ICALP)*. 2018, 27:1–27:16. ISBN: 978-3-95977-076-7. DOI: [10.4230/LIPIcs.ICALP.2018.27](https://doi.org/10.4230/LIPIcs.ICALP.2018.27) (cit. on p. 69).
- [CKLM20] Mahdi Cheraghchi, Valentine Kabanets, Zhenjian Lu, and Dimitrios Myrisiotis. “Circuit lower bounds for MCSP from local pseudorandom generators”. In: *ACM Trans. Comput. Theory* 12.3 (2020), Art. 21, 27. ISSN: 1942-3454. DOI: [10.1145/3404860](https://doi.org/10.1145/3404860) (cit. on p. 106).
- [CL20] Eshan Chattopadhyay and Jyun-Jie Liao. “Optimal Error Pseudodistributions for Read-Once Branching Programs”. In: *Proc. 35th Annual IEEE Conference on Computational Complexity (CCC)*. Vol. 169. 2020, 25:1–25:27. ISBN: 978-3-95977-156-6. DOI: [10.4230/LIPIcs.CCC.2020.25](https://doi.org/10.4230/LIPIcs.CCC.2020.25) (cit. on pp. 12, 51).
- [CL21] Kuan Cheng and Xin Li. “Efficient document exchange and error correcting codes with asymmetric information”. In: *Proc. 2021 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2021, pp. 2424–2443. DOI: [10.1137/1.9781611976465.144](https://doi.org/10.1137/1.9781611976465.144) (cit. on p. 16).
- [CLK11] Timothy H. Click, Aibing Liu, and George A. Kaminski. “Quality of random number generators significantly affects results of Monte Carlo simulations for organic and biological systems”. In: *Journal of Computational Chemistry* 32.3 (2011), pp. 513–524. DOI: [10.1002/jcc.21638](https://doi.org/10.1002/jcc.21638) (cit. on p. 6).
- [CLLO21] Lijie Chen, Zhenjian Lu, Xin Lyu, and Igor C. Oliveira. “Majority vs. Approximate Linear Sum and Average-Case Complexity Below NC¹”. In: *Proc. 48th International Colloquium on Automata, Languages and Programming (ICALP)*. 2021, 51:1–51:20. ISBN: 978-3-95977-195-5. DOI: [10.4230/LIPIcs.ICALP.2021.51](https://doi.org/10.4230/LIPIcs.ICALP.2021.51) (cit. on p. 62).
- [CLW20] Lijie Chen, Xin Lyu, and R. Ryan Williams. “Almost-everywhere circuit lower bounds from non-trivial derandomization”. In: *Proc. 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2020, pp. 1–12. DOI: [10.1109/FOCS46700.2020.00009](https://doi.org/10.1109/FOCS46700.2020.00009) (cit. on p. 69).
- [CM08] Sebastian M. Cioabă and M. Ram Murty. “Expander graphs and gaps between primes”. In: *Forum Math.* 20.4 (2008), pp. 745–756. ISSN: 0933-7741. DOI: [10.1515/FORUM.2008.035](https://doi.org/10.1515/FORUM.2008.035) (cit. on p. 46).
- [CNS99] Jin-Yi Cai, Ajay Nerurkar, and D. Sivakumar. “Hardness and hierarchy theorems for probabilistic quasi-polynomial time”. In: *Proc. 31st Annual ACM Symposium on Theory of Computing (STOC)*. 1999, pp. 726–735. DOI: [10.1145/301250.301444](https://doi.org/10.1145/301250.301444) (cit. on p. 69).
- [Cod94] P.D. Coddington. “Analysis of Random Number Generators Using Monte Carlo Simulation”. In: *International Journal of Modern Physics C* 05.03 (1994), pp. 547–560. DOI: [10.1142/S0129183194000726](https://doi.org/10.1142/S0129183194000726) (cit. on p. 6).
- [Coh16] Michael B. Cohen. “Ramanujan graphs in polynomial time”. In: *Proc. 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2016, pp. 276–281. DOI: [10.1109/FOCS.2016.37](https://doi.org/10.1109/FOCS.2016.37) (cit. on p. 46).
- [CR22] Lijie Chen and Hanlin Ren. “Strong average-case circuit lower bounds from nontrivial derandomization”. In: *SIAM J. Comput.* 51.3 (2022), STOC20–115–STOC20–173. ISSN: 0097-5397. DOI: [10.1137/20M1364886](https://doi.org/10.1137/20M1364886) (cit. on p. 69).

- [CRS00] Suresh Chari, Pankaj Rohatgi, and Aravind Srinivasan. “Improved algorithms via approximations of probability distributions”. In: *J. Comput. System Sci.* 61.1 (2000), pp. 81–107. ISSN: 0022-0000. DOI: [10.1006/jcss.1999.1695](https://doi.org/10.1006/jcss.1999.1695) (cit. on p. 32).
- [CRSW13] L. Elisa Celis, Omer Reingold, Gil Segev, and Udi Wieder. “Balls and bins: smaller hash families and faster evaluation”. In: *SIAM J. Comput.* 42.3 (2013), pp. 1030–1050. ISSN: 0097-5397. DOI: [10.1137/120871626](https://doi.org/10.1137/120871626) (cit. on p. 98).
- [CRTY20] Lijie Chen, Ron D. Rothblum, Roei Tell, and Eylon Yogev. “On Exponential-Time Hypotheses, Derandomization, and Circuit Lower Bounds: Extended Abstract”. In: *Proc. 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2020, pp. 13–23. DOI: [10.1109/FOCS46700.2020.00010](https://doi.org/10.1109/FOCS46700.2020.00010) (cit. on p. 69).
- [CT21] Lijie Chen and Roei Tell. “Simple and Fast Derandomization from Very Hard Functions: Eliminating Randomness at Almost No Cost”. In: *Proc. 53rd Annual ACM Symposium on Theory of Computing (STOC)*. 2021, 283–291. ISBN: 9781450380539. DOI: [10.1145/3406325.3451059](https://doi.org/10.1145/3406325.3451059) (cit. on pp. 68, 69).
- [De11] Anindya De. “Pseudorandomness for Permutation and Regular Branching Programs”. In: *Proc. 26th Annual IEEE Conference on Computational Complexity (CCC)*. 2011, pp. 221–231. DOI: [10.1109/CCC.2011.23](https://doi.org/10.1109/CCC.2011.23) (cit. on pp. 55, 109).
- [DETT10] Anindya De, Omid Etesami, Luca Trevisan, and Madhur Tulsiani. “Improved pseudorandom generators for depth 2 circuits”. In: *Proc. 14th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 2010, pp. 504–517. DOI: [10.1007/978-3-642-15369-3_38](https://doi.org/10.1007/978-3-642-15369-3_38) (cit. on pp. 32, 35, 90).
- [DGJSV10] Ilias Diakonikolas, Parikshit Gopalan, Ragesh Jaiswal, Rocco A. Servedio, and Emanuele Viola. “Bounded independence fools halfspaces”. In: *SIAM J. Comput.* 39.8 (2010), pp. 3441–3462. ISSN: 0097-5397. DOI: [10.1137/100783030](https://doi.org/10.1137/100783030) (cit. on p. 9).
- [DHH19] Dean Doron, Pooya Hatami, and William M. Hoza. “Near-Optimal Pseudorandom Generators for Constant-Depth Read-Once Formulas”. In: *Proc. 34th Computational Complexity Conference (CCC)*. 2019, 16:1–16:34. ISBN: 978-3-95977-116-0. DOI: [10.4230/LIPIcs.CCC.2019.16](https://doi.org/10.4230/LIPIcs.CCC.2019.16) (cit. on pp. 102, 108).
- [DHH20] Dean Doron, Pooya Hatami, and William M. Hoza. “Log-Seed Pseudorandom Generators via Iterated Restrictions”. In: *Proc. 35th Computational Complexity Conference (CCC)*. 2020, 6:1–6:36. ISBN: 978-3-95977-156-6. DOI: [10.4230/LIPIcs.CCC.2020.6](https://doi.org/10.4230/LIPIcs.CCC.2020.6) (cit. on pp. 43, 100, 102, 108).
- [DMOZ20] Dean Doron, Dana Moshkovitz, Justin Oh, and David Zuckerman. “Nearly Optimal Pseudorandomness From Hardness”. In: *Proc. 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2020, pp. 1057–1068. DOI: [10.1109/FOCS46700.2020.00102](https://doi.org/10.1109/FOCS46700.2020.00102) (cit. on pp. 7, 68, 69).
- [DMRTV21] Dean Doron, Raghu Meka, Omer Reingold, Avishay Tal, and Salil Vadhan. “Pseudorandom Generators for Read-Once Monotone Branching Programs”. In: *Proc. 25th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 2021, 58:1–58:21. ISBN: 978-3-95977-207-5. DOI: [10.4230/LIPIcs.APPROX/RANDOM.2021.58](https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2021.58) (cit. on pp. 98, 100, 102, 108).
- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. “Fuzzy extractors: how to generate strong keys from biometrics and other noisy data”. In: *SIAM J. Comput.* 38.1 (2008), pp. 97–139. ISSN: 0097-5397. DOI: [10.1137/060651380](https://doi.org/10.1137/060651380) (cit. on pp. 57, 58).

- [DVV16] Akshay Degwekar, Vinod Vaikuntanathan, and Prashant Nalini Vasudevan. “Fine-grained cryptography”. In: *Proc. 36th Annual International Cryptology Conference (CRYPTO)*. 2016, pp. 533–562. DOI: [10.1007/978-3-662-53015-3_19](https://doi.org/10.1007/978-3-662-53015-3_19) (cit. on p. 69).
- [EFF85] P. Erdős, P. Frankl, and Z. Füredi. “Families of finite sets in which no set is covered by the union of r others”. In: *Israel J. Math.* 51.1-2 (1985), pp. 79–89. ISSN: 0021-2172. DOI: [10.1007/BF02772959](https://doi.org/10.1007/BF02772959) (cit. on p. 66).
- [EGLNV98] Guy Even, Oded Goldreich, Michael Luby, Noam Nisan, and Boban Veličković. “Efficient approximation of product distributions”. In: *Random Structures Algorithms* 13.1 (1998), pp. 1–16. ISSN: 1042-9832. DOI: [10.1002/\(SICI\)1098-2418\(199808\)13:1<1::AID-RSA1>3.0.CO;2-W](https://doi.org/10.1002/(SICI)1098-2418(199808)13:1<1::AID-RSA1>3.0.CO;2-W) (cit. on p. 47).
- [EWT21] Shohei Egashira, Yuyu Wang, and Keisuke Tanaka. “Fine-grained cryptography revisited”. In: *J. Cryptology* 34.3 (2021), Paper No. 23, 43. ISSN: 0933-2790. DOI: [10.1007/s00145-021-09390-3](https://doi.org/10.1007/s00145-021-09390-3) (cit. on p. 69).
- [Fil10] Yuval Filmus. *Smolensky’s Lower Bound*. 2010. URL: <https://yuvalfilmus.cs.technion.ac.il/Manuscripts/Smolensky.pdf> (cit. on p. 69).
- [FK18] Michael A. Forbes and Zander Kelley. “Pseudorandom generators for read-once branching programs, in any order”. In: *Proc. 59th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2018, pp. 946–955. DOI: [10.1109/FOCS.2018.00093](https://doi.org/10.1109/FOCS.2018.00093) (cit. on pp. 9, 86, 92, 94, 96, 97, 108, 109).
- [FLW92] Alan M. Ferrenberg, D. P. Landau, and Y. Joanna Wong. “Monte Carlo simulations: Hidden errors from ‘good’ random number generators”. In: *Phys. Rev. Lett.* 69 (23 1992), pp. 3382–3384. DOI: [10.1103/PhysRevLett.69.3382](https://doi.org/10.1103/PhysRevLett.69.3382) (cit. on p. 6).
- [FMF85] Thomas Filk, Mihail Marcu, and Klaus Fredenhagen. “Long range correlations in random number generators and their influence on Monte Carlo simulations”. In: *Physics Letters B* 165.1 (1985), pp. 125–130. ISSN: 0370-2693. DOI: [10.1016/0370-2693\(85\)90705-1](https://doi.org/10.1016/0370-2693(85)90705-1) (cit. on p. 6).
- [Fri08] Joel Friedman. “A proof of Alon’s second eigenvalue conjecture and related problems”. In: *Mem. Amer. Math. Soc.* 195.910 (2008), pp. viii+100. ISSN: 0065-9266. DOI: [10.1090/memo/0910](https://doi.org/10.1090/memo/0910) (cit. on p. 46).
- [Fri93] Joel Friedman. “Some geometric aspects of graphs and their eigenfunctions”. In: *Duke Math. J.* 69.3 (1993), pp. 487–525. ISSN: 0012-7094. DOI: [10.1215/S0012-7094-93-06921-9](https://doi.org/10.1215/S0012-7094-93-06921-9) (cit. on p. 46).
- [FSUV13] Bill Fefferman, Ronen Shaltiel, Christopher Umans, and Emanuele Viola. “On Beating the Hybrid Argument”. In: *Theory Comput.* 9 (2013), pp. 809–843. DOI: [10.4086/toc.2013.v009a026](https://doi.org/10.4086/toc.2013.v009a026) (cit. on p. 69).
- [Für88] Zoltán Füredi. “Matchings and covers in hypergraphs”. In: *Graphs Combin.* 4.2 (1988), pp. 115–206. ISSN: 0911-0119. DOI: [10.1007/BF01864160](https://doi.org/10.1007/BF01864160) (cit. on p. 66).
- [GKL93] Oded Goldreich, Hugo Krawczyk, and Michael Luby. “On the existence of pseudorandom generators”. In: *SIAM J. Comput.* 22.6 (1993), pp. 1163–1175. ISSN: 0097-5397. DOI: [10.1137/0222069](https://doi.org/10.1137/0222069) (cit. on p. 69).
- [GKM18] Parikshit Gopalan, Daniel M. Kane, and Raghu Meka. “Pseudorandomness via the discrete Fourier transform”. In: *SIAM J. Comput.* 47.6 (2018), pp. 2451–2487. ISSN: 0097-5397. DOI: [10.1137/16M1062132](https://doi.org/10.1137/16M1062132) (cit. on p. 9).

- [GL89] O. Goldreich and L. A. Levin. “A Hard-Core Predicate for All One-Way Functions”. In: *Proc. 21st Annual ACM Symposium on Theory of Computing (STOC)*. 1989, pp. 25–32. ISBN: 0897913078. DOI: [10.1145/73007.73010](https://doi.org/10.1145/73007.73010) (cit. on p. 69).
- [GMRTV12] Parikshit Gopalan, Raghu Meka, Omer Reingold, Luca Trevisan, and Salil Vadhan. “Better Pseudorandom Generators from Milder Pseudorandom Restrictions”. In: *Proc. 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2012, pp. 120–129. DOI: [10.1109/FOCS.2012.77](https://doi.org/10.1109/FOCS.2012.77) (cit. on pp. 25, 47, 51, 93, 100, 102).
- [GMT82] Shafi Goldwasser, Silvio Micali, and Po Tong. “Why and how to establish a private code on a public network”. In: *Proc. 23rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 1982, pp. 134–144. DOI: [10.1109/SFCS.1982.100](https://doi.org/10.1109/SFCS.1982.100) (cit. on p. 69).
- [Gol01] Oded Goldreich. *Foundations of Cryptography Volume I: Basic Tools*. Cambridge University Press, 2001. DOI: [10.1017/CB09780511546891](https://doi.org/10.1017/CB09780511546891) (cit. on p. 69).
- [Gol10] Oded Goldreich. *A primer on pseudorandom generators*. Vol. 55. University Lecture Series. American Mathematical Society, Providence, RI, 2010, pp. x+114. ISBN: 978-0-8218-5192-0. DOI: [10.1090/ulect/055](https://doi.org/10.1090/ulect/055) (cit. on p. 9).
- [Gol11] Oded Goldreich. “In a world of $\mathbf{P} = \mathbf{BPP}$ ”. In: *Studies in complexity and cryptography*. Vol. 6650. Lecture Notes in Comput. Sci. Springer, Heidelberg, 2011, pp. 191–232. DOI: [10.1007/978-3-642-22670-0_20](https://doi.org/10.1007/978-3-642-22670-0_20) (cit. on p. 69).
- [Gra93] Peter Grassberger. “On correlations in ‘good’ random number generators”. In: *Physics Letters A* 181.1 (1993), pp. 43–46. ISSN: 0375-9601. DOI: [10.1016/0375-9601\(93\)91122-L](https://doi.org/10.1016/0375-9601(93)91122-L) (cit. on p. 6).
- [GUV09] Venkatesan Guruswami, Christopher Umans, and Salil Vadhan. “Unbalanced expanders and randomness extractors from Parvaresh-Vardy codes”. In: *J. ACM* 56.4 (2009), Art. 20, 34. ISSN: 0004-5411. DOI: [10.1145/1538902.1538904](https://doi.org/10.1145/1538902.1538904) (cit. on pp. 57, 59).
- [GVW11] Oded Goldreich, Salil Vadhan, and Avi Wigderson. “Simplified derandomization of BPP using a hitting set generator”. In: *Studies in Complexity and Cryptography*. Vol. 6650. Lecture Notes in Computer Science. Springer, Heidelberg, 2011, pp. 59–67. DOI: [10.1007/978-3-642-22670-0_8](https://doi.org/10.1007/978-3-642-22670-0_8) (cit. on p. 12).
- [GY20] Parikshit Gopalan and Amir Yehudayoff. “Concentration for limited independence via inequalities for the elementary symmetric polynomials”. In: *Theory Comput.* 16 (2020), Paper No. 17, 29. DOI: [10.4086/toc.2020.v016a017](https://doi.org/10.4086/toc.2020.v016a017) (cit. on pp. 47, 110).
- [Hås01] Johan Håstad. “A slight sharpening of LMN”. In: *Journal of Computer and System Sciences* 63.3 (2001), pp. 498–508. DOI: [10.1006/jcss.2001.1803](https://doi.org/10.1006/jcss.2001.1803) (cit. on p. 38).
- [Hås14] Johan Håstad. “On the Correlation of Parity and Small-Depth Circuits”. In: *SIAM J. Comput.* 43.5 (2014), pp. 1699–1708. DOI: [10.1137/120897432](https://doi.org/10.1137/120897432) (cit. on pp. 63, 72).
- [Has89] John Hastad. “Almost Optimal Lower Bounds for Small Depth Circuits”. In: *Adv. Comput. Res.* 5 (1989), pp. 143–170. URL: <https://www.csc.kth.se/~johanh/largesmalldepth.pdf> (cit. on pp. 68, 72).

- [HHR06] Iftach Haitner, Danny Harnik, and Omer Reingold. “Efficient pseudorandom generators from exponentially hard one-way functions”. In: *Proc. 33rd International Colloquium on Automata, Languages and Programming (ICALP)*. 2006, pp. 228–239. DOI: [10.1007/11787006_20](https://doi.org/10.1007/11787006_20) (cit. on p. 69).
- [HHR11] Iftach Haitner, Danny Harnik, and Omer Reingold. “On the Power of the Randomized Iterate”. In: *SIAM J. Comput.* 40.6 (2011), pp. 1486–1528. ISSN: 0097-5397. DOI: [10.1137/080721820](https://doi.org/10.1137/080721820) (cit. on p. 69).
- [HHTT22] Pooya Hatami, William M. Hoza, Avishay Tal, and Roei Tell. “Fooling constant-depth threshold circuits”. In: *Proc. 62nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 2022, pp. 104–115. DOI: [10.1109/FOCS52979.2021.00019](https://doi.org/10.1109/FOCS52979.2021.00019) (cit. on pp. 43, 106, 108, 109).
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. “A pseudo-random generator from any one-way function”. In: *SIAM J. Comput.* 28.4 (1999), pp. 1364–1396. ISSN: 0097-5397. DOI: [10.1137/S0097539793244708](https://doi.org/10.1137/S0097539793244708) (cit. on pp. 7, 69).
- [HLV18] Elad Haramaty, Chin Ho Lee, and Emanuele Viola. “Bounded independence plus noise fools products”. In: *SIAM J. Comput.* 47.2 (2018), pp. 493–523. ISSN: 0097-5397. DOI: [10.1137/17M1129088](https://doi.org/10.1137/17M1129088) (cit. on pp. 47, 93).
- [Hol06] Thomas Holenstein. “Pseudorandom generators from one-way functions: a simple construction for any hardness”. In: *Theory of cryptography*. Vol. 3876. Lecture Notes in Comput. Sci. Springer, Berlin, 2006, pp. 443–461. DOI: [10.1007/11681878_23](https://doi.org/10.1007/11681878_23) (cit. on p. 69).
- [Hoz21] William M. Hoza. “Better pseudodistributions and derandomization for space-bounded computation”. In: *Proc. 25th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 2021, 28:1–28:23. DOI: [10.4230/LIPIcs.APPROX/RANDOM.2021.28](https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2021.28) (cit. on pp. 8, 12, 51, 60).
- [Hoz22] William M. Hoza. “Recent Progress on Derandomizing Space-Bounded Computation”. In: *Bulletin of the EATCS* 138 (2022), pp. 114–143. URL: <https://eatcs.org/images/bulletin/beatcs138.pdf> (cit. on p. 49).
- [HPV21] William M. Hoza, Edward Pyne, and Salil Vadhan. “Pseudorandom Generators for Unbounded-Width Permutation Branching Programs”. In: *Proc. 12th Conference on Innovations in Theoretical Computer Science (ITCS)*. 2021, 7:1–7:20. ISBN: 978-3-95977-177-1. DOI: [10.4230/LIPIcs.ITCS.2021.7](https://doi.org/10.4230/LIPIcs.ITCS.2021.7) (cit. on pp. 12, 55).
- [HR03] Tzvikia Hartman and Ran Raz. “On the distribution of the number of roots of polynomials and explicit weak designs”. In: *Random Structures Algorithms* 23.3 (2003), pp. 235–263. ISSN: 1042-9832. DOI: [10.1002/rsa.10095](https://doi.org/10.1002/rsa.10095) (cit. on pp. 67, 68).
- [HRV13] Iftach Haitner, Omer Reingold, and Salil Vadhan. “Efficiency improvements in constructing pseudorandom generators from one-way functions”. In: *SIAM J. Comput.* 42.3 (2013), pp. 1405–1430. ISSN: 0097-5397. DOI: [10.1137/100814421](https://doi.org/10.1137/100814421) (cit. on p. 69).
- [HS19] Prahladh Harsha and Srikanth Srinivasan. “On polynomial approximations to AC^0 ”. In: *Random Structures Algorithms* 54.2 (2019), pp. 289–303. DOI: [10.1002/rsa.20786](https://doi.org/10.1002/rsa.20786) (cit. on pp. 37, 38, 41, 42).

- [HZ20] William M. Hoza and David Zuckerman. “Simple optimal hitting sets for small-success \mathbf{RL} ”. In: *SIAM J. Comput.* 49.4 (2020), pp. 811–820. ISSN: 0097-5397. DOI: [10.1137/19M1268707](https://doi.org/10.1137/19M1268707) (cit. on pp. 25, 51, 60).
- [IMP12] Russell Impagliazzo, William Matthews, and Ramamohan Paturi. “A satisfiability algorithm for AC^0 ”. In: *Proc. 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2012, pp. 961–972. URL: <https://dl.acm.org/doi/10.5555/2095116.2095193> (cit. on p. 63).
- [IMZ19] Russell Impagliazzo, Raghu Meka, and David Zuckerman. “Pseudorandomness from shrinkage”. In: *J. ACM* 66.2 (2019), Art. 11, 16. ISSN: 0004-5411. DOI: [10.1145/3230630](https://doi.org/10.1145/3230630) (cit. on pp. 102, 105, 106).
- [INW94] Russell Impagliazzo, Noam Nisan, and Avi Wigderson. “Pseudorandomness for Network Algorithms”. In: *Proc. 26th Annual ACM Symposium on Theory of Computing (STOC)*. 1994, pp. 356–364. ISBN: 0897916638. DOI: [10.1145/195058.195190](https://doi.org/10.1145/195058.195190) (cit. on pp. 9, 45, 48, 50, 68, 109, 110).
- [ISW06] Russell Impagliazzo, Ronen Shaltiel, and Avi Wigderson. “Reducing the seed length in the Nisan-Wigderson generator”. In: *Combinatorica* 26.6 (2006), pp. 647–681. ISSN: 0209-9683. DOI: [10.1007/s00493-006-0036-8](https://doi.org/10.1007/s00493-006-0036-8) (cit. on pp. 68, 69).
- [ISW99] Russell Impagliazzo, Ronen Shaltiel, and Avi Wigderson. “Near-optimal conversion of hardness into pseudo-randomness”. In: *Proc. 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 1999, pp. 181–190. DOI: [10.1109/SFFCS.1999.814590](https://doi.org/10.1109/SFFCS.1999.814590) (cit. on p. 12).
- [IW01] Russell Impagliazzo and Avi Wigderson. “Randomness vs time: derandomization under a uniform assumption”. In: *J. Comput. System Sci.* 63.4 (2001), pp. 672–688. ISSN: 0022-0000. DOI: [10.1006/jcss.2001.1780](https://doi.org/10.1006/jcss.2001.1780) (cit. on p. 69).
- [IW97] Russell Impagliazzo and Avi Wigderson. “P = BPP If E Requires Exponential Circuits: Derandomizing the XOR Lemma”. In: *Proc. 29th Annual ACM Symposium on Theory of Computing (STOC)*. 1997, pp. 220–229. ISBN: 0897918886. DOI: [10.1145/258533.258590](https://doi.org/10.1145/258533.258590) (cit. on pp. 7, 11, 69).
- [KC00] Valentine Kabanets and Jin-Yi Cai. “Circuit Minimization Problem”. In: *Proc. 32nd Annual ACM Symposium on Theory of Computing (STOC)*. 2000, 73–79. DOI: [10.1145/335305.335314](https://doi.org/10.1145/335305.335314) (cit. on p. 10).
- [Kel21] Zander Kelley. “An Improved Derandomization of the Switching Lemma”. In: *Proc. 53rd Annual ACM Symposium on Theory of Computing (STOC)*. 2021, 272–282. ISBN: 9781450380539. DOI: [10.1145/3406325.3451054](https://doi.org/10.1145/3406325.3451054) (cit. on pp. 89, 90).
- [KL18] Valentine Kabanets and Zhenjian Lu. “Satisfiability and Derandomization for Small Polynomial Threshold Circuits”. In: *Proc. 22nd International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 2018, 46:1–46:19. ISBN: 978-3-95977-085-9. DOI: [10.4230/LIPIcs.APPROX-RANDOM.2018.46](https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2018.46) (cit. on p. 67).
- [KLW10] Adam R. Klivans, Homin K. Lee, and Andrew Wan. “Mansour’s Conjecture is True for Random DNF Formulas”. In: *Proc. 23rd Conference on Learning Theory (COLT)*. 2010, pp. 368–380. URL: <http://www.learningtheory.org/colt2010/papers/085Lee.pdf> (cit. on p. 34).

- [KM02] Adam R. Klivans and Dieter van Melkebeek. “Graph Nonisomorphism Has Subexponential Size Proofs Unless the Polynomial-Time Hierarchy Collapses”. In: *SIAM J. Comput.* 31.5 (2002), pp. 1501–1526. ISSN: 0097-5397. DOI: [10.1137/S0097539700389652](https://doi.org/10.1137/S0097539700389652) (cit. on pp. 7, 67).
- [KM93] Eyal Kushilevitz and Yishay Mansour. “Learning decision trees using the Fourier spectrum”. In: *SIAM J. Comput.* 22.6 (1993), pp. 1331–1348. ISSN: 0097-5397. DOI: [10.1137/0222080](https://doi.org/10.1137/0222080) (cit. on pp. 24, 109).
- [KNP11] Michal Koucký, Prajakta Nimbhorkar, and Pavel Pudlák. “Pseudorandom Generators for Group Products”. In: *Proc. 43rd Annual ACM Symposium on Theory of Computing (STOC)*. 2011, pp. 263–272. DOI: [10.1145/1993636.1993672](https://doi.org/10.1145/1993636.1993672) (cit. on pp. 55, 109).
- [KNW08] Daniel M. Kane, Jelani Nelson, and David P. Woodruff. *Revisiting Norm Estimation in Data Streams*. 2008. arXiv: [0811.3648 \[cs.DS\]](https://arxiv.org/abs/0811.3648) (cit. on pp. 57, 59, 109).
- [KW84] Claus Kalle and Stephan Wansleben. “Problems with the random number generator RANF implemented on the CDC cyber 205”. In: *Computer Physics Communications* 33.4 (1984), pp. 343–346. ISSN: 0010-4655. DOI: [10.1016/0010-4655\(84\)90139-5](https://doi.org/10.1016/0010-4655(84)90139-5) (cit. on p. 6).
- [Lee19] Chin Ho Lee. “Fourier Bounds and Pseudorandom Generators for Product Tests”. In: *Proc. 34th Computational Complexity Conference (CCC)*. 2019, 7:1–7:25. DOI: [10.4230/LIPIcs.CCC.2019.7](https://doi.org/10.4230/LIPIcs.CCC.2019.7) (cit. on pp. 47, 100, 102).
- [Lev87] Leonid A. Levin. “One way functions and pseudorandom generators”. In: *Combinatorica* 7.4 (1987), pp. 357–363. ISSN: 0209-9683. DOI: [10.1007/BF02579323](https://doi.org/10.1007/BF02579323) (cit. on p. 69).
- [LLSZ97] Nathan Linial, Michael Luby, Michael Saks, and David Zuckerman. “Efficient construction of a small hitting set for combinatorial rectangles in high dimension”. In: *Combinatorica* 17.2 (1997), pp. 215–234. ISSN: 0209-9683. DOI: [10.1007/BF01200907](https://doi.org/10.1007/BF01200907) (cit. on p. 47).
- [LMN93] Nathan Linial, Yishay Mansour, and Noam Nisan. “Constant depth circuits, Fourier transform, and learnability”. In: *Journal of the ACM* 40.3 (1993), pp. 607–620. DOI: [10.1145/174130.174138](https://doi.org/10.1145/174130.174138) (cit. on p. 38).
- [LN90] Nathan Linial and Noam Nisan. “Approximate inclusion-exclusion”. In: *Combinatorica* 10.4 (1990), pp. 349–365. DOI: [10.1007/BF02128670](https://doi.org/10.1007/BF02128670) (cit. on p. 37).
- [Lov09] Shachar Lovett. “Unconditional Pseudorandom Generators for Low Degree Polynomials”. In: *Theory of Computing* 5.1 (2009), pp. 69–82. DOI: [10.4086/toc.2009.v005a003](https://doi.org/10.4086/toc.2009.v005a003) (cit. on p. 26).
- [LPS88] A. Lubotzky, R. Phillips, and P. Sarnak. “Ramanujan graphs”. In: *Combinatorica* 8.3 (1988), pp. 261–277. ISSN: 0209-9683. DOI: [10.1007/BF02126799](https://doi.org/10.1007/BF02126799) (cit. on p. 46).
- [LPV22] Chin Ho Lee, Edward Pyne, and Salil Vadhan. “Fourier Growth of Regular Branching Programs”. In: *Proc. 26th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 2022, 2:1–2:21. ISBN: 978-3-95977-249-5. DOI: [10.4230/LIPIcs.APPROX/RANDOM.2022.2](https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2022.2) (cit. on pp. 55, 81, 83, 86).

- [LS07] Pierre L’Ecuyer and Richard Simard. “TestU01: A C Library for Empirical Testing of Random Number Generators”. In: *ACM Trans. Math. Softw.* 33.4 (Aug. 2007). ISSN: 0098-3500. DOI: [10.1145/1268776.1268777](https://doi.org/10.1145/1268776.1268777) (cit. on p. 6).
- [LS11] Shachar Lovett and Srikanth Srinivasan. “Correlation bounds for poly-size AC^0 circuits with $n^{1-o(1)}$ symmetric gates”. In: *Proc. 15th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 2011, pp. 640–651. DOI: [10.1007/978-3-642-22935-0_54](https://doi.org/10.1007/978-3-642-22935-0_54) (cit. on p. 67).
- [Lu02] Chi-Jen Lu. “Improved pseudorandom generators for combinatorial rectangles”. In: *Combinatorica* 22.3 (2002), pp. 417–433. ISSN: 0209-9683. DOI: [10.1007/s004930200021](https://doi.org/10.1007/s004930200021) (cit. on p. 47).
- [LV17] Chin Ho Lee and Emanuele Viola. “Some limitations of the sum of small-bias distributions”. In: *Theory Comput.* 13 (2017), Paper No. 16, 23. DOI: [10.4086/toc.2017.v013a016](https://doi.org/10.4086/toc.2017.v013a016) (cit. on pp. 32, 51).
- [LV20] Chin Ho Lee and Emanuele Viola. “More on bounded independence plus noise: pseudorandom generators for read-once polynomials”. In: *Theory Comput.* 16 (2020), Paper No. 7, 50. DOI: [10.4086/toc.2020.v016a007](https://doi.org/10.4086/toc.2020.v016a007) (cit. on p. 102).
- [LV96] M. Luby and B. Veličković. “On Deterministic Approximation of DNF”. In: *Algorithmica* 16.4/5 (1996), pp. 415–433. DOI: [10.1007/BF01940873](https://doi.org/10.1007/BF01940873) (cit. on p. 44).
- [LVW93] Michael Luby, Boban Veličković, and Avi Wigderson. “Deterministic approximate counting of depth-2 circuits”. In: *Proc. 2nd Israel Symposium on Theory and Computing Systems (ISTCS)*. 1993, pp. 18–24. DOI: [10.1109/ISTCS.1993.253488](https://doi.org/10.1109/ISTCS.1993.253488) (cit. on pp. 26, 67).
- [LW06] Michael Luby and Avi Wigderson. “Pairwise Independence and Derandomization”. In: *Foundations and Trends in Theoretical Computer Science* 1.4 (2006), pp. 237–301. ISSN: 1551-305X. DOI: [10.1561/0400000009](https://doi.org/10.1561/0400000009) (cit. on p. 9).
- [Lyu22] Xin Lyu. “Improved Pseudorandom Generators for AC^0 Circuits”. In: *Proc. 37th Computational Complexity Conference (CCC)*. 2022, 34:1–34:25. ISBN: 978-3-95977-241-9. DOI: [10.4230/LIPIcs.CCC.2022.34](https://doi.org/10.4230/LIPIcs.CCC.2022.34) (cit. on pp. 34, 89, 90, 102, 108).
- [Mar88] G. A. Margulis. “Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators”. In: *Problemy Peredachi Informatsii* 24.1 (1988), pp. 51–60. ISSN: 0555-2923. URL: <http://mi.mathnet.ru/eng/ppi/v24/i1/p51> (cit. on p. 46).
- [MBH86] A Milchev, K Binder, and DW Heermann. “Fluctuations and lack of self-averaging in the kinetics of domain growth”. In: *Zeitschrift für Physik B Condensed Matter* 63.4 (1986), pp. 521–535. DOI: [10.1007/BF01726202](https://doi.org/10.1007/BF01726202) (cit. on p. 6).
- [Mil01] Peter Bro Miltersen. “Derandomizing complexity classes”. In: *Handbook of randomized computing, Vol. I, II*. Vol. 9. Comb. Optim. Kluwer Acad. Publ., Dordrecht, 2001, pp. 843–941. DOI: [10.1007/978-1-4615-0013-1_19](https://doi.org/10.1007/978-1-4615-0013-1_19) (cit. on p. 9).
- [MN98] Makoto Matsumoto and Takuji Nishimura. “Mersenne Twister: A 623-Dimensionally Equidistributed Uniform Pseudo-Random Number Generator”. In: *ACM Trans. Model. Comput. Simul.* 8.1 (Jan. 1998), 3–30. ISSN: 1049-3301. DOI: [10.1145/272991.272995](https://doi.org/10.1145/272991.272995) (cit. on p. 6).

- [MOP20] Sidhanth Mohanty, Ryan O’Donnell, and Pedro Paredes. “Explicit near-Ramanujan graphs of every degree”. In: *Proc. 52nd Annual ACM Symposium on Theory of Computing (STOC)*. ACM, New York, 2020, pp. 510–523. DOI: [10.1145/3357713.3384231](https://doi.org/10.1145/3357713.3384231) (cit. on p. 46).
- [Mor94] Moshe Morgenstern. “Existence and explicit constructions of $q+1$ regular Ramanujan graphs for every prime power q ”. In: *J. Combin. Theory Ser. B* 62.1 (1994), pp. 44–62. ISSN: 0095-8956. DOI: [10.1006/jctb.1994.1054](https://doi.org/10.1006/jctb.1994.1054) (cit. on p. 46).
- [MRT19] Raghu Meka, Omer Reingold, and Avishay Tal. “Pseudorandom Generators for Width-3 Branching Programs”. In: *Proc. 51st Annual ACM Symposium on Theory of Computing (STOC)*. 2019, pp. 626–637. DOI: [10.1145/3313276.3316319](https://doi.org/10.1145/3313276.3316319) (cit. on pp. 43, 50, 56, 102, 109).
- [MSS15] Adam W. Marcus, Daniel A. Spielman, and Nikhil Srivastava. “Interlacing families I: Bipartite Ramanujan graphs of all degrees”. In: *Ann. of Math. (2)* 182.1 (2015), pp. 307–325. ISSN: 0003-486X. DOI: [10.4007/annals.2015.182.1.7](https://doi.org/10.4007/annals.2015.182.1.7) (cit. on p. 46).
- [MZ09] Raghu Meka and David Zuckerman. “Small-bias spaces for group products”. In: *Proc. 13th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 2009, pp. 658–672. DOI: [10.1007/978-3-642-03685-9_49](https://doi.org/10.1007/978-3-642-03685-9_49) (cit. on p. 51).
- [MZ13] Raghu Meka and David Zuckerman. “Pseudorandom generators for polynomial threshold functions”. In: *SIAM J. Comput.* 42.3 (2013), pp. 1275–1301. ISSN: 0097-5397. DOI: [10.1137/100811623](https://doi.org/10.1137/100811623) (cit. on p. 9).
- [MZ21] Noam Mazon and Jiapeng Zhang. “Simple Constructions from (Almost) Regular One-Way Functions”. In: *Proc. 19th Theory of Cryptography Conference (TCC)*. 2021, pp. 457–485. ISBN: 978-3-030-90453-1. DOI: [10.1007/978-3-030-90453-1_16](https://doi.org/10.1007/978-3-030-90453-1_16) (cit. on p. 69).
- [Nil91] A. Nilli. “On the second eigenvalue of a graph”. In: *Discrete Math.* 91.2 (1991), pp. 207–210. ISSN: 0012-365X. DOI: [10.1016/0012-365X\(91\)90112-F](https://doi.org/10.1016/0012-365X(91)90112-F) (cit. on p. 46).
- [Nis91] N. Nisan. “Pseudorandom bits for constant depth circuits”. In: *Combinatorica* 11.1 (1991), pp. 63–70. DOI: [10.1007/BF01375474](https://doi.org/10.1007/BF01375474) (cit. on p. 67).
- [Nis92] Noam Nisan. “Pseudorandom generators for space-bounded computation”. In: *Combinatorica* 12.4 (1992), pp. 449–461. DOI: [10.1007/BF01305237](https://doi.org/10.1007/BF01305237) (cit. on pp. 50, 68, 109).
- [NN93] Joseph Naor and Moni Naor. “Small-bias probability spaces: efficient constructions and applications”. In: *SIAM J. Comput.* 22.4 (1993), pp. 838–856. ISSN: 0097-5397. DOI: [10.1137/0222053](https://doi.org/10.1137/0222053) (cit. on pp. 19, 21, 110).
- [NW94] Noam Nisan and Avi Wigderson. “Hardness vs Randomness”. In: *J. Comput. Syst. Sci.* 49.2 (1994), pp. 149–167. DOI: [10.1016/S0022-0000\(05\)80043-1](https://doi.org/10.1016/S0022-0000(05)80043-1) (cit. on pp. 7, 9, 64, 67).
- [NZ96] Noam Nisan and David Zuckerman. “Randomness is linear in space”. In: *J. Comput. System Sci.* 52.1 (1996), pp. 43–52. ISSN: 0022-0000. DOI: [10.1006/jcss.1996.0004](https://doi.org/10.1006/jcss.1996.0004) (cit. on pp. 56, 59).
- [O’D14] Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014. DOI: [10.1017/CB09781139814782](https://doi.org/10.1017/CB09781139814782) (cit. on p. 9).

- [Per13] Nicole Perlroth. “Government Announces Steps to Restore Confidence on Encryption Standards”. In: *The New York Times* (2013). URL: <https://bits.blogs.nytimes.com/2013/09/10/government-announces-steps-to-restore-confidence-on-encryption-standards/> (visited on 07/14/2021) (cit. on p. 7).
- [Per90] Rene Peralta. “On the randomness complexity of algorithms”. In: *University of Wisconsin, Milwaukee CS Research Report TR 90-1* (1990) (cit. on p. 19).
- [PR85] Giorgio Parisi and Federico Rapuano. “Effects of the random number generator on computer simulations”. In: *Physics Letters B* 157.4 (1985), pp. 301–302. ISSN: 0370-2693. DOI: [10.1016/0370-2693\(85\)90670-7](https://doi.org/10.1016/0370-2693(85)90670-7) (cit. on p. 6).
- [PS84] Christos H. Papadimitriou and Michael Sipser. “Communication complexity”. In: *J. Comput. System Sci.* 28.2 (1984), pp. 260–269. ISSN: 0022-0000. DOI: [10.1016/0022-0000\(84\)90069-2](https://doi.org/10.1016/0022-0000(84)90069-2) (cit. on p. 103).
- [PV21a] Edward Pyne and Salil Vadhan. “Limitations of the Impagliazzo-Nisan-Wigderson Pseudorandom Generator Against Permutation Branching Programs”. In: *Proc. 27th International Computing and Combinatorics Conference (COCOON)*. 2021, pp. 3–12. DOI: [10.1007/978-3-030-89543-3_1](https://doi.org/10.1007/978-3-030-89543-3_1) (cit. on pp. 55, 56).
- [PV21b] Edward Pyne and Salil Vadhan. “Pseudodistributions that beat all pseudorandom generators (extended abstract)”. In: *Proc. 36th Computational Complexity Conference (CCC)*. 2021, 33:1–33:15. DOI: [10.4230/LIPIcs.CCC.2021.33](https://doi.org/10.4230/LIPIcs.CCC.2021.33). Full version: ECCC preprint [TR21-019](https://eccc.weizmann.ac.il/report/2021/019) (cit. on pp. 12, 51, 55).
- [PV22] Edward Pyne and Salil Vadhan. “Deterministic Approximation of Random Walks via Queries in Graphs of Unbounded Size”. In: *Proc. 5th Symposium on Simplicity in Algorithms (SOSA)*. 2022, pp. 57–67. DOI: [10.1137/1.9781611977066.5](https://doi.org/10.1137/1.9781611977066.5) (cit. on pp. 12, 55).
- [Raz09] Alexander Razborov. “A Simple Proof of Bazzi’s Theorem”. In: *ACM Transactions on Computation Theory* 1.1 (2009). DOI: [10.1145/1490270.1490273](https://doi.org/10.1145/1490270.1490273) (cit. on pp. 37, 39).
- [Raz87] A. A. Razborov. “Lower bounds on the size of bounded depth circuits over a complete basis with logical addition”. In: *Math. Notes* 41.4 (1987), pp. 333–338. DOI: [10.1007/BF01137685](https://doi.org/10.1007/BF01137685) (cit. on p. 41).
- [Raz91] Alexander A. Razborov. “Lower bounds for deterministic and nondeterministic branching programs”. In: *Proc. 8th International Conference on Fundamentals of Computation Theory (FCT)*. 1991, pp. 47–60. DOI: [10.1007/3-540-54458-5_49](https://doi.org/10.1007/3-540-54458-5_49) (cit. on p. 49).
- [Röd85] Vojtěch Rödl. “On a Packing and Covering Problem”. In: *European Journal of Combinatorics* 6.1 (1985), pp. 69–78. ISSN: 0195-6698. DOI: [10.1016/S0195-6698\(85\)80023-8](https://doi.org/10.1016/S0195-6698(85)80023-8) (cit. on p. 67).
- [Ros19] Benjamin Rossman. “Criticality of Regular Formulas”. In: *Proc. 34th Computational Complexity Conference (CCC)*. 2019, 1:1–1:28. ISBN: 978-3-95977-116-0. DOI: [10.4230/LIPIcs.CCC.2019.1](https://doi.org/10.4230/LIPIcs.CCC.2019.1) (cit. on pp. 72, 89).
- [RRV02] Ran Raz, Omer Reingold, and Salil Vadhan. “Extracting all the randomness and reducing the error in Trevisan’s extractors”. In: *J. Comput. System Sci.* 65.1 (2002), pp. 97–128. ISSN: 0022-0000. DOI: [10.1006/jcss.2002.1824](https://doi.org/10.1006/jcss.2002.1824) (cit. on p. 68).

- [RS10] Y. Rabani and A. Shpilka. “Explicit Construction of a Small ε -Net for Linear Threshold Functions”. In: *SIAM J. on Computing* 39.8 (2010), pp. 3501–3520. DOI: [10.1137/090764190](https://doi.org/10.1137/090764190) (cit. on p. 9).
- [RSV13] Omer Reingold, Thomas Steinke, and Salil Vadhan. “Pseudorandomness for regular branching programs via Fourier analysis”. In: *Proc. 17th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 2013, pp. 655–670. DOI: [10.1007/978-3-642-40328-6_45](https://doi.org/10.1007/978-3-642-40328-6_45) (cit. on pp. 53, 55, 81, 86, 87).
- [RTV06] Omer Reingold, Luca Trevisan, and Salil Vadhan. “Pseudorandom walks on regular digraphs and the **RL** vs. **L** problem”. In: *Proc. 38th Annual ACM Symposium on Theory of Computing (STOC)*. 2006, pp. 457–466. DOI: [10.1145/1132516.1132583](https://doi.org/10.1145/1132516.1132583) (cit. on p. 51).
- [RY20] Anup Rao and Amir Yehudayoff. *Communication Complexity and Applications*. Cambridge University Press, 2020. DOI: [10.1017/9781108671644](https://doi.org/10.1017/9781108671644) (cit. on p. 63).
- [Sch64] J. Schönheim. “On coverings”. In: *Pacific J. Math.* 14 (1964), pp. 1405–1411. ISSN: 0030-8730. URL: <http://projecteuclid.org/euclid.pjm/1103033815> (cit. on p. 66).
- [Sha83] Adi Shamir. “On the Generation of Cryptographically Strong Pseudorandom Sequences”. In: *ACM Trans. Comput. Syst.* 1.1 (1983), 38–44. ISSN: 0734-2071. DOI: [10.1145/357353.357357](https://doi.org/10.1145/357353.357357) (cit. on pp. 65, 69).
- [Sko22] Maciej Skorski. “Tight Chernoff-Like Bounds Under Limited Independence”. In: *Proc. 26th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 2022, 15:1–15:14. ISBN: 978-3-95977-249-5. DOI: [10.4230/LIPIcs.APPROX/RANDOM.2022.15](https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2022.15) (cit. on p. 104).
- [Smo87] Roman Smolensky. “Algebraic Methods in the Theory of Lower Bounds for Boolean Circuit Complexity”. In: *Proc. 19th Annual ACM Symposium on Theory of Computing (STOC)*. 1987, pp. 77–82. DOI: [10.1145/28395.28404](https://doi.org/10.1145/28395.28404) (cit. on p. 41).
- [Smo93] R. Smolensky. “On Representations by Low-Degree Polynomials”. In: *Proc. 34th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 1993, pp. 130–138. DOI: [10.1109/SFCS.1993.366874](https://doi.org/10.1109/SFCS.1993.366874) (cit. on p. 69).
- [ST18] Rocco A. Servedio and Li-Yang Tan. “Luby-Veličković-Wigderson revisited: improved correlation bounds and pseudorandom generators for depth-two circuits”. In: *Proc. 22nd International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 2018, 56:1–56:20. DOI: [10.4230/LIPIcs.APPROX-RANDOM.2018.56](https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2018.56) (cit. on p. 67).
- [ST19a] Rocco A. Servedio and Li-Yang Tan. “Improved Pseudorandom Generators from Pseudorandom Multi-Switching Lemmas”. In: *Proc. 28th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*. 2019, 45:1–45:23. ISBN: 978-3-95977-125-2. DOI: [10.4230/LIPIcs.APPROX-RANDOM.2019.45](https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2019.45) (cit. on pp. 88, 90).
- [ST19b] Rocco A. Servedio and Li-Yang Tan. “Pseudorandomness for read- k DNF formulas”. In: *Proc. 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2019, pp. 621–638. DOI: [10.1137/1.9781611975482.39](https://doi.org/10.1137/1.9781611975482.39) (cit. on p. 34).

- [Ste12] Thomas Steinke. *Pseudorandomness for Permutation Branching Programs Without the Group Theory*. ECCC preprint TR12-083. 2012. URL: <https://eccc.weizmann.ac.il/report/2012/083/> (cit. on pp. 55, 109).
- [STV01] Madhu Sudan, Luca Trevisan, and Salil Vadhan. “Pseudorandom Generators without the XOR Lemma”. In: *J. Comput. Syst. Sci.* 62.2 (2001), pp. 236–266. DOI: [10.1006/jcss.2000.1730](https://doi.org/10.1006/jcss.2000.1730) (cit. on p. 69).
- [SU05] Ronen Shaltiel and Christopher Umans. “Simple extractors for all min-entropies and a new pseudorandom generator.” In: *J. ACM* 52.2 (2005), pp. 172–216. DOI: [10.1145/1059513.1059516](https://doi.org/10.1145/1059513.1059516) (cit. on pp. 68, 69).
- [Sub61] B. A. Subbotovskaya. “Realizations of linear function by formulas using $+$, \cdot , $-$ ”. In: *Doklady Akademii Nauk SSSR* 136:3 (1961). In Russian, pp. 553–555. URL: <http://mi.mathnet.ru/eng/dan/v136/i3/p553> (cit. on p. 71).
- [SVW17] Thomas Steinke, Salil Vadhan, and Andrew Wan. “Pseudorandomness and Fourier-Growth Bounds for Width-3 Branching Programs”. In: *Theory Comput.* 13 (2017), Paper No. 12. DOI: [10.4086/toc.2017.v013a012](https://doi.org/10.4086/toc.2017.v013a012) (cit. on pp. 87, 98).
- [ŠŽ21] Jiří Šíma and Stanislav Žák. “A polynomial-time construction of a hitting set for read-once branching programs of width 3”. In: *Fund. Inform.* 184.4 (2021), pp. 307–354. ISSN: 0169-2968. DOI: [10.3233/fi-2021-2101](https://doi.org/10.3233/fi-2021-2101) (cit. on p. 51).
- [SZ99] Michael Saks and Shiyu Zhou. “ $\text{BP}_H\text{SPACE}(S) \subseteq \text{DSPACE}(S^{3/2})$ ”. In: *J. Comput. System Sci.* 58.2 (1999), pp. 376–403. ISSN: 0022-0000. DOI: [10.1006/jcss.1998.1616](https://doi.org/10.1006/jcss.1998.1616) (cit. on p. 8).
- [Ta-15] Amnon Ta-Shma. *Randomized algorithms and de-randomization*. Course at Tel-Aviv University. 2015. URL: <http://www.cs.tau.ac.il/~amnon/Classes/2015-PRG/class.htm> (cit. on p. 9).
- [Ta-16] Amnon Ta-Shma. *Expanders, Pseudorandomness and Derandomization*. Course at Tel-Aviv University. 2016. URL: <http://www.cs.tau.ac.il/~amnon/Classes/2016-PRG/class.htm> (cit. on p. 9).
- [Ta-18] Amnon Ta-Shma. *Space-Bounded Computation*. Course at Tel-Aviv University. 2018. URL: <http://www.cs.tau.ac.il/~amnon/Classes/2018-Space/class.htm> (cit. on p. 9).
- [Ta-19] Amnon Ta-Shma. *A first course in derandomization*. Course at Tel-Aviv University. 2019. URL: <http://www.cs.tau.ac.il/~amnon/Classes/2019-Derandomization/class.htm> (cit. on p. 9).
- [Tal17] Avishay Tal. “Tight Bounds on the Fourier Spectrum of AC^0 ”. In: *Proc. 32nd Computational Complexity Conference (CCC)*. 2017, 15:1–15:31. ISBN: 978-3-95977-040-8. DOI: [10.4230/LIPIcs.CCC.2017.15](https://doi.org/10.4230/LIPIcs.CCC.2017.15) (cit. on pp. 37, 38, 81, 88, 90).
- [Tal21] Avishay Tal. *Pseudorandomness*. Course at University of California, Berkeley. 2021. URL: <https://www.avishaytal.org/pseudorandomness> (cit. on p. 9).
- [Tar93] Jun Tarui. “Probabilistic polynomials, AC^0 functions and the polynomial-time hierarchy”. In: *Theoret. Comput. Sci.* 113.1 (1993), pp. 167–183. ISSN: 0304-3975. DOI: [10.1016/0304-3975\(93\)90214-E](https://doi.org/10.1016/0304-3975(93)90214-E) (cit. on pp. 41, 42).
- [TO92] Seinosuke Toda and Mitsunori Ogiwara. “Counting classes are at least as hard as the polynomial-time hierarchy”. In: *SIAM J. Comput.* 21.2 (1992), pp. 316–328. ISSN: 0097-5397. DOI: [10.1137/0221023](https://doi.org/10.1137/0221023) (cit. on p. 41).

- [Tre01] Luca Trevisan. “Extractors and pseudorandom generators”. In: *J. ACM* 48.4 (2001), pp. 860–879. ISSN: 0004-5411. DOI: [10.1145/502090.502099](https://doi.org/10.1145/502090.502099) (cit. on p. 67).
- [Tre05] Luca Trevisan. *Pseudorandomness and Combinatorial Constructions*. Course at University of California, Berkeley. 2005. URL: <https://web.archive.org/web/20150115081847/http://www.cs.berkeley.edu/~luca/pacc/> (cit. on p. 9).
- [TS17] Amnon Ta-Shma. “Explicit, almost optimal, epsilon-balanced codes”. In: *Proc. 49th Annual ACM Symposium on Theory of Computing (STOC)*. 2017, pp. 238–251. DOI: [10.1145/3055399.3055408](https://doi.org/10.1145/3055399.3055408) (cit. on p. 20).
- [TV07] Luca Trevisan and Salil Vadhan. “Pseudorandomness and average-case complexity via uniform reductions”. In: *Comput. Complexity* 16.4 (2007), pp. 331–364. ISSN: 1016-3328. DOI: [10.1007/s00037-007-0233-x](https://doi.org/10.1007/s00037-007-0233-x) (cit. on p. 69).
- [TX13] Luca Trevisan and Tongke Xue. “A Derandomized Switching Lemma and an Improved Derandomization of AC⁰”. In: *Proc. 28th Annual IEEE Conference on Computational Complexity (CCC)*. 2013, pp. 242–247. DOI: [10.1109/CCC.2013.32](https://doi.org/10.1109/CCC.2013.32) (cit. on pp. 88, 90).
- [Uma03] Christopher Umans. “Pseudo-random generators for all hardnesses”. In: *J. of Computer and System Sciences* 67.2 (2003), pp. 419–440. DOI: [10.1016/S0022-0000\(03\)00046-1](https://doi.org/10.1016/S0022-0000(03)00046-1) (cit. on pp. 7, 68, 69).
- [Vad12] Salil P. Vadhan. “Pseudorandomness”. In: *Foundations and Trends in Theoretical Computer Science* 7.1-3 (2012), pp. 1–336. DOI: [10.1561/0400000010](https://doi.org/10.1561/0400000010) (cit. on pp. 9, 45, 46, 57, 58).
- [Vig17] Sebastiano Vigna. “Further scramblings of Marsaglia’s xorshift generators”. In: *J. Comput. Appl. Math.* 315 (2017), pp. 175–181. ISSN: 0377-0427. DOI: [10.1016/j.cam.2016.11.006](https://doi.org/10.1016/j.cam.2016.11.006) (cit. on p. 6).
- [Vio07] Emanuele Viola. “Pseudorandom Bits for Constant-Depth Circuits with Few Arbitrary Symmetric Gates”. In: *SIAM J. Comput.* 36.5 (2007), pp. 1387–1403. DOI: [10.1137/050640941](https://doi.org/10.1137/050640941) (cit. on p. 67).
- [Vio09] Emanuele Viola. “The Sum of D Small-Bias Generators Fools Polynomials of Degree D ”. In: *Comput. Complexity* 18.2 (2009), pp. 209–217. DOI: [10.1007/s00037-009-0273-5](https://doi.org/10.1007/s00037-009-0273-5) (cit. on pp. 9, 26, 62, 110).
- [Vio14] Emanuele Viola. “Randomness buys depth for approximate counting”. In: *Comput. Complexity* 23.3 (2014), pp. 479–508. ISSN: 1016-3328. DOI: [10.1007/s00037-013-0076-6](https://doi.org/10.1007/s00037-013-0076-6) (cit. on p. 47).
- [Vio17] Emanuele Viola. *Special Topics in Complexity Theory*. Course at Northeastern University. 2017. URL: <https://www.ccs.neu.edu/home/viola/classes/spepf17.html> (cit. on p. 9).
- [Vio22] Emanuele Viola. *Correlation bounds against polynomials*. ECCC preprint TR22-142. 2022. URL: <https://eccc.weizmann.ac.il/report/2022/142/> (cit. on p. 29).
- [VV86] L. G. Valiant and V. V. Vazirani. “NP is as easy as detecting unique solutions”. In: *Theoret. Comput. Sci.* 47.1 (1986), pp. 85–93. ISSN: 0304-3975. DOI: [10.1016/0304-3975\(86\)90135-0](https://doi.org/10.1016/0304-3975(86)90135-0) (cit. on p. 42).

- [VZ12] Salil Vadhan and Colin Jia Zheng. “Characterizing pseudoentropy and simplifying pseudorandom generator constructions”. In: *Proc. 44th Annual ACM Symposium on Theory of Computing (STOC)*. 2012, pp. 817–836. DOI: [10.1145/2213977.2214051](https://doi.org/10.1145/2213977.2214051) (cit. on p. 69).
- [Weg87] Ingo Wegener. *The complexity of Boolean functions*. Wiley-Teubner Series in Computer Science. John Wiley & Sons, Inc., 1987, pp. xii+457. ISBN: 0-471-91555-6. URL: <https://dl.acm.org/doi/10.5555/35517> (cit. on p. 49).
- [Yao82] Andrew C. Yao. “Theory and Applications of Trapdoor Functions”. In: *Proc. 23rd Annual ACM Symposium on Theory of Computing (STOC)*. 1982, pp. 80–91. DOI: [10.1109/SFCS.1982.45](https://doi.org/10.1109/SFCS.1982.45) (cit. on pp. 7, 69).
- [YGLW15] Yu Yu, Dawu Gu, Xiangxue Li, and Jian Weng. “The randomized iterate, revisited—almost linear seed length PRGs from a broader class of one-way functions”. In: *Proc. 12th Theory of Cryptography Conference (TCC)*. 2015, pp. 7–35. DOI: [10.1007/978-3-662-46494-6_2](https://doi.org/10.1007/978-3-662-46494-6_2) (cit. on p. 69).
- [YLW15] Yu Yu, Xiangxue Li, and Jian Weng. “Pseudorandom generators from regular one-way functions: new constructions with improved parameters”. In: *Theoret. Comput. Sci.* 569 (2015), pp. 58–69. ISSN: 0304-3975. DOI: [10.1016/j.tcs.2014.12.013](https://doi.org/10.1016/j.tcs.2014.12.013) (cit. on p. 69).
- [Zuc01] David Zuckerman. *Pseudorandomness and Combinatorial Constructions*. Course at the University of Texas at Austin. 2001. URL: <https://www.cs.utexas.edu/~diz/395T/01/> (cit. on p. 9).
- [Zuc90] David Zuckerman. “General weak random sources”. In: *Proc. 31st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 1990, pp. 534–543. DOI: [10.1109/FSCS.1990.89574](https://doi.org/10.1109/FSCS.1990.89574) (cit. on p. 56).