

Pseudorandomness

Salil P. Vadhan¹

¹ *Harvard University Cambridge, MA02138, USA, salil@seas.harvard.edu*

Abstract

This is a survey of *pseudorandomness*, the theory of efficiently generating objects that “look random” despite being constructed using little or no randomness. This theory has significance for a number of areas in computer science and mathematics, including computational complexity, algorithms, cryptography, combinatorics, communications, and additive number theory. Our treatment places particular emphasis on the intimate connections that have been discovered between a variety of fundamental “pseudorandom objects” that at first seem very different in nature: expander graphs, randomness extractors, list-decodable error-correcting codes, samplers, and pseudorandom generators. The structure of the presentation is meant to be suitable for teaching in a graduate-level course, with exercises accompanying each chapter.

To Kaya and Amari

Acknowledgments

My exploration of pseudorandomness began in my graduate and postdoctoral years at MIT and IAS, under the wonderful guidance of Oded Goldreich, Shafi Goldwasser, Madhu Sudan, and Avi Wigderson. It was initiated by an exciting reading group organized at MIT by Luca Trevisan, which immersed me in the subject and started my extensive collaboration with Luca. Through fortuitous circumstances, I also began to work with Omer Reingold, starting what I hope will be a lifelong collaboration. I am indebted to Oded, Shafi, Madhu, Avi, Luca, and Omer for all the insights and research experiences they have shared with me.

I have also learned a great deal from my other collaborators on pseudorandomness, including Boaz Barak, Eli Ben-Sasson, Michael Capalbo, Kai-Min Chung, Nenad Dedic, Yevgeniy Dodis, Parikshit Gopalan, Dan Gutfreund, Venkat Guruswami, Iftach Haitner, Alex Healy, Thomas Holenstein, Jesse Kamp, Danny Lewin, Adriana López-Alt, Shachar Lovett, Chi-Jen Lu, Raghu Meka, Ilya Mironov, Michael Mitzenmacher, Shien Jin Ong, Michael Rabin, Anup Rao, Ran Raz, Yakir Reshef, Leo Reyzin, Thomas Ristenpart, Eyal Rozenman, Thomas Steinke, Madhur Tulsiani, Chris Umans, Emanuele Viola, Hoeteck Wee, Colin Jia Zheng, and David Zuckerman. Needless to say, this list omits many other researchers in the field with whom I have had stimulating discussions.

The starting point for this survey was scribe notes taken by students in the 2004 version of my graduate course on Pseudorandomness. I thank those students for their contribution: Alexandr Andoni, Adi Akavia, Yan-Cheng Chang, Denis Chebikin, Hamilton Chong, Vitaly Feldman, Brian Greenberg, Chun-Yun Hsiao, Andrei Jorza, Adam Kirsch, Kevin Matulef, Mihai Pătraşcu, John Provine, Pavlo Pylyavskyy, Arthur Rudolph, Saurabh Sanghvi, Grant Schoenebeck, Jordanna Schutz, Sasha Schwartz, David Troiano, Vinod Vaikuntanathan, Kartik Venkatram, David Woodruff. I also thank the students from the other offerings of the course; Dan Gutfreund, who gave some guest lectures in 2007; and all of my teaching fellows, Johnny Chen, Kai-Min Chung, Minh Nguyen, Emanuele Viola, and Colin Jia Zheng. Special thanks are due to Levent Alpoge, Michael Forbes, Dieter van Melkebeek, Greg Price, and Adam Sealfon for their extensive feedback on the lecture notes and/or drafts of this survey.

Helpful input of various types (corrections, comments, answering questions) has also been given by Zachary Abel, Dana Albrecht, Nir Avni, Pablo Azar, Trevor Bass, Osbert Bastiani, Jeremy

Booher, Fan Chung, Ben Dozier, Chinmoy Dutta, Zhou Fan, Oded Goldreich, Andrey Grinshpun, Venkat Guruswami, Alex Healy, Stephan Holzer, Andrei Jorza, Michael von Korff, Kevn Lee, Alex Lubotzky, Avner May, Eric Miles, Shira Mitchell, Jelani Nelson, Omer Reingold, Yakir Reshef, Shubhangi Saraf, Shrenik Shah, Madhu Sudan, Justin Thaler, Jon Ullman, Ameya Velingker, Neal Wadhwa, Hoeteck Wee, Avi Wigderson, David Wu. Thanks to all of them, as well as those I have forgotten.

I am extremely grateful to James Finlay of now publishers for his years of patience and continual encouragement, without which I would have never finished this survey.

My time working on this survey included sabbatical visits to the Miller Institute for Basic Research in Science at UC Berkeley, Microsoft Research Silicon Valley, and Stanford University, as well as support from a Sloan Fellowship, a Guggenheim Fellowship, NSF grants CCF-0133096 and CCF-1116616, ONR grant N00014-04-1-0478, and US-Israel BSF grants 2006060 and 2010196.

Contents

1	Introduction	1
1.1	Overview of this Survey	1
1.2	Background Required and Teaching Tips	4
1.3	Notational Conventions	5
1.4	Chapter Notes and References	5
2	The Power of Randomness	7
2.1	Polynomial Identity Testing	7
2.2	The Computational Model and Complexity Classes	11
2.3	Sampling and Approximation Problems	16
2.4	Random Walks and S-T CONNECTIVITY	22
2.5	Exercises	27
2.6	Chapter Notes and References	31
3	Basic Derandomization Techniques	34
3.1	Enumeration	34
3.2	Nonconstructive/Nonuniform Derandomization	36
3.3	Nondeterminism	38
3.4	The Method of Conditional Expectations	39
3.5	Pairwise Independence	42
3.6	Exercises	49
3.7	Chapter Notes and References	52
4	Expander Graphs	54

4.1	Measures of Expansion	54
4.2	Random Walks on Expanders	61
4.3	Explicit Constructions	68
4.4	UNDIRECTED S-T CONNECTIVITY in Deterministic Logspace	78
4.5	Exercises	80
4.6	Chapter Notes and References	84
5	List-Decodable Codes	87
5.1	Definitions and Existence	87
5.2	List-Decoding Algorithms	93
5.3	List-decoding views of samplers and expanders	100
5.4	Expanders from Parvaresh–Vardy Codes	103
5.5	Exercises	105
5.6	Chapter Notes and References	108
6	Randomness Extractors	110
6.1	Motivation and Definition	110
6.2	Connections to Other Pseudorandom Objects	117
6.3	Constructing Extractors	124
6.4	Exercises	133
6.5	Chapter Notes and References	137
7	Pseudorandom Generators	140
7.1	Motivation and Definition	140
7.2	Cryptographic PRGs	144
7.3	Hybrid Arguments	148
7.4	Pseudorandom Generators from Average-Case Hardness	151
7.5	Worst-Case/Average-Case Reductions and Locally Decodable Codes	158
7.6	Local List Decoding and PRGs from Worst-Case Hardness	166
7.7	Connections to Other Pseudorandom Objects	172
7.8	Exercises	178
7.9	Chapter Notes and References	183
8	Conclusions	187
8.1	A Unified Theory of Pseudorandomness	187
8.2	Other Topics in Pseudorandomness	191
	References	204

1

Introduction

1.1 Overview of this Survey

Over the past few decades, *randomization* has become one of the most pervasive paradigms in computer science. Its widespread uses include:

Algorithm Design: For a number of important algorithmic problems, the most efficient algorithms known are randomized. For example:

- **PRIMALITY TESTING.** This was shown to have a randomized polynomial-time algorithm in 1977. It wasn't until 2002 that a deterministic polynomial-time algorithm was discovered. (We will see this algorithm, but not its proof.)
- **APPROXIMATE COUNTING.** Many approximate counting problems (e.g. counting perfect matchings in a bipartite graph) have randomized polynomial-time algorithms, but the fastest known deterministic algorithms take exponential time.
- **UNDIRECTED S-T CONNECTIVITY.** This was shown to have a randomized logspace algorithm in 1979. It wasn't until 2005 that a deterministic logspace algorithm was discovered — using tools from the theory of pseudorandomness, as we will see.
- **PERFECT MATCHING.** This was shown to have a randomized *polylogarithmic*-time parallel algorithm in the late 1970's. Deterministically, we only know polynomial-time algorithms.

Even in the cases where deterministic algorithms of comparable complexity were eventually found, the randomized algorithms remain considerably simpler and more efficient than the deterministic ones.

Cryptography: Randomization is central to cryptography. Indeed, cryptography is concerned with protecting secrets, and how can something be secret if it is deterministically fixed? For example, we assume that cryptographic keys are chosen at random (e.g. uniformly from the set of n -bit strings). In addition to the keys, it is known that often the cryptographic algorithms themselves

(e.g. for encryption) must be randomized to achieve satisfactory notions of security (e.g. that no partial information about the message is leaked).

Combinatorial Constructions: Randomness is often used to prove the *existence* of combinatorial objects with a desired property. Specifically, if one can show that a randomly chosen object has the property with nonzero probability, then it follows that such an object must, in fact, exist. A famous example due to Erdős is the existence of *Ramsey graphs*: A randomly chosen n -vertex graph has no clique or independent set of size $2 \log n$ with high probability. We will see several other applications of this “Probabilistic Method” in this survey, such as with two important objects mentioned below: expander graphs and error-correcting codes.

Though these *applications* of randomness are interesting and rich topics of study in their own right, they are not the focus of this survey. Rather, we ask the following:

Main Question: Can we reduce or even eliminate the use of randomness in these settings?

We have several motivations for doing this.

- **Complexity Theory:** We are interested in understanding and comparing the power of various kinds of computational resources. Since randomness is such a widely used resource, we want to know how it relates to other resources such as time, space, and parallelism. In particular, we ask: *Can every randomized algorithm be derandomized with only a small loss in efficiency?*
- **Using Physical Random Sources:** It is unclear whether the real world has physical sources of perfect randomness. We may use sources that seem to have some unpredictability, like the low order bits of a system clock or thermal noise, but these sources will generally have biases and, more problematically, correlations. Thus we ask: What can we do with a source of biased and correlated bits?
- **Explicit Constructions:** Probabilistic constructions of combinatorial objects often do not provide us with efficient algorithms for using those objects. Indeed, the randomly chosen object often has a description that is exponential in the relevant parameters. Thus, we look for *explicit* constructions — ones that are deterministic and efficient. In addition to their applications, improvements in explicit constructions serve as a measure of our progress in understanding the objects at hand. Indeed, Erdős posed the explicit construction of near-optimal Ramsey graphs as an open problem, and substantial progress on this problem was recently made using the theory of pseudorandomness (namely randomness extractors).
- **Unexpected Applications:** In addition, the theory of pseudorandomness has turned out to have many applications to problems that seem to have no connection to derandomization. These include data structures, distributed computation, circuit lower bounds and completeness results in complexity theory, reducing interaction in interactive protocols, saving memory in streaming algorithms, and more. We will see some of these applications in this survey (especially the exercises).

The paradigm we will use to study the Main Question is that of *pseudorandomness*: efficiently generating objects that “look random” using little or no randomness.

Specifically, we will study four “pseudorandom” objects:

Pseudorandom Generators: A pseudorandom generator is an algorithm that takes as input a short, perfectly random *seed* and then returns a (much longer) sequence of bits that “looks random.” That the bits output cannot be perfectly random is clear — the output is determined by the seed and there are far fewer seeds than possible bit sequences. Nevertheless, it is possible for the output to “look random” in a very meaningful and general-purpose sense. Specifically, we will require that no *efficient* algorithm can distinguish the output from a truly random sequence. The study of pseudorandom generators meeting this strong requirement originated in cryptography, where they have numerous applications. In this survey, we will emphasize their role in derandomizing algorithms.

Note that asserting that a function is a pseudorandom generator is a statement about something that efficient algorithms can’t do (in this case, distinguish two sequences). But proving that efficient algorithms cannot compute things is typically out of reach for current techniques in theoretical computer science; indeed this is why the \mathbf{P} vs. \mathbf{NP} question is so hard. Thus, we will settle for conditional statements. An ideal theorem would be something like: “If $\mathbf{P} \neq \mathbf{NP}$, then pseudorandom generators exist.” (The assumptions we make won’t exactly be $\mathbf{P} \neq \mathbf{NP}$, but hypotheses of a similar flavor.)

Randomness Extractors: A randomness extractor takes as input a source of biased and correlated bits, and then produces a sequence of almost-uniform bits as output. Their original motivation was the simulation of randomized algorithms with sources of biased and correlated bits, but they have found numerous other applications in theoretical computer science. Ideally, extractors would be deterministic, but as we will see this proves to be impossible for general sources of biased and correlated bits. Nevertheless, we will get close — constructing extractors that are only “mildly” probabilistic, in that they use small (logarithmic) number of truly random bits as a seed for the extraction.

Expander Graphs: Expanders are graphs with two seemingly contradictory properties: they are sparse (e.g. having degree that is a constant, independent of the number of vertices), but also “well-connected” in some precise sense. For example, the graph cannot be bisected without cutting a large (say, constant) fraction of the edges.

Expander graphs have numerous applications in theoretical computer science. They were originally studied for their use in designing fault-tolerant networks (e.g. for telephone lines), ones that maintain good connectivity even when links or nodes fail. But they also have less obvious applications, such as an $O(\log n)$ -time algorithm for sorting in parallel.

It is not obvious that expander graphs exist, but in fact it can be shown, via the Probabilistic Method, that a random graph of degree 3 is a “good” expander with high probability. However, many applications of expander graphs need *explicit constructions*, and these have taken longer to find. We will see some explicit constructions in this survey, but even the state-of-the-art does not always match the bounds given by the probabilistic method (in terms of the degree/expansion tradeoff).

Error-Correcting Codes: Error-correcting codes are tools for communicating over noisy channels. Specifically, they specify a way to encode messages into longer, redundant codewords so that even if the codeword gets somewhat corrupted along the way, it is still possible for the receiver to decode the original message. In his landmark paper that introduced the field of coding theory, Shannon also proved the existence of good error-correcting codes via the Probabilistic Method. That is, a random mapping of n -bit messages to $O(n)$ -bit codewords is a “good” error-correcting code with high probability. Unfortunately, these probabilistic codes are not feasible to actually use — a random mapping requires an exponentially long description, and we know of no way to decode such a mapping efficiently. Again, explicit constructions are needed.

In this survey, we will focus on the problem of *list decoding*. Specifically, we will consider scenarios where the number of corruptions is so large that unique decoding is impossible; at best one can produce a short list that is guaranteed to contain the correct message.

A Unified Theory: Each of the above objects has been the center of a large and beautiful body of research, but until recently these corpora were largely distinct. An exciting development over the past decade has been the realization that all four of these objects are almost *the same* when interpreted appropriately. Their intimate connections will be a major focus of this survey, tying together the variety of constructions and applications that we will see.

The surprise and beauty of these connections has to do with the seemingly different nature of each of these objects. Pseudorandom generators, by asserting what efficient algorithms cannot do, are objects of complexity theory. Extractors, with their focus on extracting the entropy in a correlated and biased sequence, are information-theoretic objects. Expander graphs are of course combinatorial objects (as defined above), though they can also be interpreted algebraically, as we will see. Error-correcting codes involve a mix of combinatorics, information theory, and algebra. Because of the connections, we obtain new perspectives on each of the objects, and make substantial advances on our understanding of each by translating intuitions and techniques from the study of the others.

1.2 Background Required and Teaching Tips

The presentation assumes a good undergraduate background in the theory of computation, and general mathematical maturity. Specifically, it is assumed that the reader is familiar with basic algorithms and discrete mathematics, e.g. as covered in [CLRS], including some exposure to randomized algorithms; and with basic computational complexity including P, NP, and reductions, e.g. as covered in [Sip3]. Experience with elementary abstract algebra, particularly finite fields, is helpful; recommended texts are [Art, LN].

Most of the material in the survey is covered in a one-semester graduate course that the author teaches at Harvard University, which consists of 24 lectures of 1.5 hours each. Most of the students in that course take at least one graduate-level course in theoretical computer science before this one.

The exercises are an important part of the survey, as they include proofs of some key facts, introduce some concepts that will be used in later chapters, and illustrate applications of the material to other topics. Problems that are particularly challenging or require more creativity than most are marked with a star.

1.3 Notational Conventions

We denote the set of numbers $\{1, \dots, n\}$ by $[n]$. We write \mathbb{N} for the set of nonnegative integers (so we consider 0 to be a natural number). We write $S \subset T$ to mean that S is a subset of T (not necessarily strict), and $S \subsetneq T$ for S being a strict subset of T . An inequality we use often is $\binom{n}{k} \leq (ne/k)^k$. All logarithms are base 2 unless otherwise specified. We often use the convention that lowercase letters are the logarithm (base 2) of the corresponding capital letter (e.g. $N = 2^n$).

Throughout, we consider random variables that can take values in arbitrary discrete sets (as well as real-valued random variables). We generally use capital letters, e.g. X , to denote random variables and lowercase letters, e.g. x , to denote specific values. We write $x \stackrel{R}{\leftarrow} X$ to indicate that x is sampled according to X . For a set S , we write $x \stackrel{R}{\leftarrow} S$ to mean that x is selected uniformly at random from S . We use the convention that multiple occurrences of a random variable in an expression refer to the same instantiation, e.g. $\Pr[X = X] = 1$. The *support* of a random variable X is denote $\text{Supp}(X) \stackrel{\text{def}}{=} \{x : \Pr[X = x] > 0\}$. For an event E , we write $X|_E$ to denote the random variable X conditioned on the event E . For a set S , we write U_S to denote a random variable distributed uniformly over S . For $n \in \mathbb{N}$, U_n is a random variable distributed uniformly over $\{0, 1\}^n$.

1.4 Chapter Notes and References

In this chapter, we only provide pointers to general surveys and textbooks on the topics covered, plus citations for specific results mentioned.

General introductions to the theory of pseudorandomness (other than this survey) include [Gol8, Mil2, Tre5].

Recommended textbooks focused on randomized algorithms are [MU, MR]. The specific randomized and deterministic algorithms mentioned are due to [Mil1, Rab, SS, AKS, KLM, JSV, AKL⁺, Rei2, Lov1, KUW]; for more details see Section 2.6.

Recommended textbooks on cryptography are [Gol3, Gol4, KL]. The idea that encryption should be randomized is due to Goldwasser and Micali [GM].

The Probabilistic Method for combinatorial constructions is the subject of the book [AS1]. Erdős used this method to prove the existence of Ramsey graphs in [Erd1]. Major recent progress on explicit constructions of Ramsey graphs was recently obtained by Barak, Rao, Shaltiel, and Wigderson [BRSW] via the theory of randomness extractors.

The modern notion of a pseudorandom generator was formulated in the works of Blum and Micali [BM] and Yao [Yao], motivated by cryptographic applications. We will spend most of our time on a variant of the Blum–Micali–Yao notion, proposed by Nisan and Wigderson [NW], where the generator is allowed more running time than the algorithms it fools. A detailed treatment of the Blum–Micali–Yao notion can be found in [Gol3].

Surveys on randomness extractors are [NT, Sha1, Sha3]. The notion of extractor that we will focus on is the one due to Nisan and Zuckerman [NZ].

A detailed survey of expander graphs is [HLW]. The probabilistic construction of expander graphs is due to Pinsker [Pin]. The application of expanders to sorting in parallel is due to Ajtai, Komlós, and Szemerédi [AKS1].

A classic text on coding theory is [MS]. For a modern, computer-science-oriented treatment, we recommend Sudan’s lecture notes [Sud3]. Shannon started this field with the paper [Sha3]. The

notion of list decoding was proposed by Elias [Eli1] and Wozencraft [Woz], and was reinvigorated in the work of Sudan [Sud1]. Recent progress on list decoding is covered in [Gur2].

2

The Power of Randomness

Before we study the derandomization of randomized algorithms, we will need some algorithms to derandomize. Thus in this chapter, we present a number of examples of randomized algorithms, as well as develop the complexity-theoretic framework and basic tools for studying randomized algorithms.

2.1 Polynomial Identity Testing

In this section, we give a randomized algorithm to solve the following computational problem.

Computational Problem 2.1. POLYNOMIAL IDENTITY TESTING: given two multivariate polynomials, $f(x_1, \dots, x_n)$ and $h(x_1, \dots, x_n)$, decide whether $f = g$.

This problem statement requires some clarification. Specifically, we need to say what we mean by:

- “polynomials”: A (*multivariate*) *polynomial* is a finite expression of the form

$$f(x_1, \dots, x_n) = \sum_{i_1, \dots, i_n \in \mathbb{N}} c_{i_1, \dots, i_n} x_1^{i_1} x_2^{i_2} \cdots x_n^{i_n}.$$

We need to specify from what space the coefficients of the polynomials come; they could be the integers, reals, rationals, etc. In general, we will assume that the coefficients are chosen from a *field* (a set with addition and multiplication, where every nonzero element has a multiplicative inverse) or more generally an (*integral*) *domain* (where the product of two nonzero elements is always nonzero). Examples of fields include \mathbb{Q} (the rationals), \mathbb{R} (the reals), and \mathbb{Z}_p (integers modulo p) where p is prime. An integral domain that is not a field is \mathbb{Z} (the integers), but every integral domain is contained in its *field of fractions*, which is \mathbb{Q} in the case of \mathbb{Z} . \mathbb{Z}_n for composite n is not even an integral domain. We remark that there does exist a finite field \mathbb{F}_q of size $q = p^k$ for every prime p and

positive integer k , and in fact this field is unique (up to isomorphism). Note that \mathbb{F}_q is only equal to \mathbb{Z}_q in case q is prime (i.e. $k = 1$). For more background on algebra, see the references in the chapter notes.

For a polynomial $f(x_1, \dots, x_n) = \sum_{i_1, \dots, i_n} c_{i_1, \dots, i_n} x_1^{i_1} x_2^{i_2} \cdots x_n^{i_n}$, we define its *degree* (a.k.a. *total degree*) to be the maximum sum of the exponents $i_1 + \cdots + i_n$ over its monomials with nonzero coefficients c_{i_1, \dots, i_n} . Its *degree in x_j* is the maximum of i_j over its monomials with nonzero coefficients.

- “ $f = g$ ”: What does it mean for two polynomials to be equal? There are two natural choices: the polynomials are the same as *functions* (they have the same output for every point in the domain), or the polynomials are the same as *formal polynomials* (the coefficients for each monomial are the same).

These two definitions are equivalent over the integers, but they are *not* equivalent over finite fields. For example, consider

$$f(x) = \prod_{\alpha \in \mathbb{F}} (x - \alpha).$$

for a finite field \mathbb{F} .¹ It is easy to see that $f(\alpha) = 0$ for all $\alpha \in \mathbb{F}$, but $f \neq 0$ as a formal polynomial. For us, equality refers to equality as formal polynomials unless otherwise specified (but often we’ll be working with domains that are larger than the degrees of the polynomials, in which case the two definitions are equivalent, as follows from the Schwartz–Zippel Lemma below).

- “given”: What does it mean to be given a polynomial? There are several possibilities here:
 - (1) As a list of coefficients: this trivializes the problem of POLYNOMIAL IDENTITY TESTING, as we can just compare. Since we measure complexity as a function of the input length, this algorithm runs in linear time.
 - (2) As an “oracle”: a black box that, given any point in the domain, gives the value of the polynomial. Here the only “explicit” inputs to the algorithm are the description of the field \mathbb{F} , the number n of variables, and a bound d on the degree of the polynomial, and we consider an algorithm to be efficient if it runs in time polynomial in the lengths of these inputs and $n \log |\mathbb{F}|$ (since it takes time at least $n \log |\mathbb{F}|$ to write down an input to the oracle).
 - (3) As an *arithmetic formula*: a sequence of symbols like $(x_1 + x_2)(x_3 + x_7 + 6x_5)x_3(x_5 - x_6) + x_2x_4(2x_3 + 3x_5)$ that describes the polynomial. Observe that while we can solve POLYNOMIAL IDENTITY TESTING by expanding the polynomials and grouping terms, the expanded polynomials may have length exponential in the length of the formula, and thus the algorithm is not efficient as a function of the input length.

More general than formulas are circuits. An *arithmetic circuit* consists of a directed acyclic graph, consisting of *input nodes*, which have indegree 0 and are labeled by input variables or constants, and *computation nodes*, which have indegree 2 and are labelled by operations ($+$ or \times) specifying how to compute a

¹When expanded and terms are collected, this polynomial p can be shown to simply equal $x^{|\mathbb{F}|} - x$.

value given the values at its children; one of the computation nodes is designated as the *output node*. Observe that every arithmetic circuit defines a polynomial in its input variables x_1, \dots, x_n . Arithmetic formulas are equivalent to arithmetic circuits where the underlying graph is a tree.

The randomized algorithm we describe will work for all of the formulations above, but is only interesting for the second and third ones (oracles and arithmetic circuits/formulas). It will be convenient to work with the following equivalent version of the problem.

Computational Problem 2.2. POLYNOMIAL IDENTITY TESTING (reformulation): Given a polynomial $f(x_1, \dots, x_n)$, is $f = 0$?

That is, we consider the special case of the original problem where $g = 0$. Any solution for the general version of course implies one for the special case; conversely, we can solve the general version by applying the special case to the polynomial $f' = f - g$.

Algorithm 2.3 (POLYNOMIAL IDENTITY TESTING).

Input: A multivariate polynomial $f(x_1, \dots, x_n)$ of degree at most d over a field/domain \mathbb{F} .

- (1) Let $S \subset \mathbb{F}$ be any set of size $2d$.
- (2) Choose $\alpha_1, \dots, \alpha_n \stackrel{\text{R}}{\leftarrow} S$.
- (3) Evaluate $f(\alpha_1, \dots, \alpha_n)$. If the result is 0, accept. Otherwise, reject.

It is clear that if $f = 0$, the algorithm will always accept. The correctness in case $f \neq 0$ is based on the following simple but very useful lemma.

Lemma 2.4 (Schwartz–Zippel Lemma). If f is a nonzero polynomial of degree d over a field (or integral domain) \mathbb{F} and $S \subset \mathbb{F}$, then

$$\Pr_{\alpha_1, \dots, \alpha_n \stackrel{\text{R}}{\leftarrow} S} [f(\alpha_1, \dots, \alpha_n) = 0] \leq \frac{d}{|S|}.$$

In the univariate case ($n = 1$), this amounts to the familiar fact that a polynomial with coefficients in a field and degree d has at most d roots. The proof for multivariate polynomials proceeds by induction on n , and we leave it as an exercise (Problem 2.1).

By the Schwartz-Zippel lemma, the algorithm will err with probability at most $1/2$ when $f \neq 0$. This error probability can be reduced by repeating the algorithm many times (or by increasing $|S|$). Note that the error probability is only over the coin tosses of the algorithm, not over the input polynomial f . This is what we mean when we say *randomized algorithm*; it should work on a worst-case input with high probability over the coin tosses of the algorithm. In contrast, algorithms whose correctness (or efficiency) only holds for inputs that are randomly chosen (according to a particular distribution) are called *heuristics*, and their study is called *average-case analysis*.

Note that we need a few things to ensure that our algorithm will work.

- First, we need a bound on the degree of the polynomial. We can get this in different ways depending on how the polynomial is represented. For example, for arithmetic formulas, the degree is bounded by the length of the formula. For arithmetic circuits, the degree is at most exponential in the size (or even depth) of the circuit.
- We also must be able to evaluate f when the variables take arbitrary values in some set S of size $2d$. For starters, this requires that the domain \mathbb{F} is of size at least $2d$. We should also have an explicit description of the domain \mathbb{F} enabling us to write down and manipulate field elements, and randomly sample from a subset of size at least $2d$ (e.g. the description of $\mathbb{F} = \mathbb{Z}_p$ is simply the prime p). Then, if we are given f as an oracle, we have the ability to evaluate f by definition. If we are given f as an arithmetic formula or circuit, then we can do a bottom-up, gate-by-gate evaluation. However, over infinite domains (like \mathbb{Z}), there is subtlety — the bit-length of the numbers can grow exponentially large. Problem 2.4 gives a method for coping with this.

Since these two conditions are satisfied, we have a polynomial-time randomized algorithm for POLYNOMIAL IDENTITY TESTING for polynomials given as arithmetic formulas over \mathbb{Z} (or even circuits, by Problem 2.4). There are no known subexponential-time *deterministic* algorithms for this problem, even for formulas in $\Sigma\Pi\Sigma$ form (i.e. a sum of terms, each of which is the product of linear functions in the input variables). A deterministic polynomial-time algorithm for $\Sigma\Pi\Sigma$ formulas where the outermost sum has only a constant number of terms was obtained quite recently (2005).

2.1.1 Application to Perfect Matching

Now we will see an application of POLYNOMIAL IDENTITY TESTING to an important graph-theoretic problem.

Definition 2.5. Let $G = (V, E)$ be a graph. A *matching* on G is a set $E' \subset E$ such that no two edges in E' have a common endpoint. A *perfect matching* is a matching such that every vertex is incident to an edge in the matching.

Computational Problem 2.6. PERFECT MATCHING: given a graph G , decide whether there is a perfect matching in G .

Unlike POLYNOMIAL IDENTITY TESTING, PERFECT MATCHING is known to have deterministic polynomial-time algorithms — e.g. using alternating paths, or by reduction to MAX FLOW in the bipartite case. However, both of these algorithms seem to be inherently sequential in nature. With randomization, we can obtain an efficient parallel algorithm.

Algorithm 2.7 (PERFECT MATCHING in bipartite graphs).

Input: a bipartite graph G with vertices numbered $1, \dots, n$ on each side and edges $E \subseteq [n] \times [n]$.²

²Recall that $[n]$ denotes the set $\{1, \dots, n\}$. See Section 1.3.

We construct an $n \times n$ matrix A where

$$A_{i,j}(x) = \begin{cases} x_{i,j} & \text{if } (i, j) \in E \\ 0 & \text{otherwise} \end{cases},$$

where $x_{i,j}$ is a formal variable.

Consider the multivariate polynomial

$$\det(A(x)) = \sum_{\sigma \in S_n} \text{sign}(\sigma) \cdot \prod_i A_{i,\sigma(i)},$$

where S_n denotes the set of permutations on $[n]$. Note that the σ 'th term is nonzero if and only if the permutation σ defines a perfect matching. That is, $(i, \sigma(i)) \in E$ for all $1 \leq i \leq n$. So $\det(A(x)) = 0$ iff G has no perfect matching. Moreover its degree is bounded by n , and given values $\alpha_{i,j}$ for the $x_{i,j}$'s we can evaluate $\det(A(\alpha))$ efficiently in parallel (in polylogarithmic time using an appropriate algorithm for the determinant).

So to test for a perfect matching efficiently in parallel, just run the POLYNOMIAL IDENTITY TESTING algorithm with, say, $S = \{1, \dots, 2n\} \subset \mathbb{Z}$, to test whether $\det(A(x)) = 0$.

Some remarks:

- The above also provides the most efficient *sequential* algorithm for PERFECT MATCHING, using the fact that DETERMINANT has the same time complexity as MATRIX MULTIPLICATION, which is known to be at most $O(n^{2.38})$.
- More sophisticated versions of the algorithm apply to nonbipartite graphs, and enable *finding* perfect matchings in the same parallel or sequential time complexity (where the result for sequential time is quite recent).
- POLYNOMIAL IDENTITY TESTING has been also used to obtain a randomized algorithm for PRIMALITY TESTING, which was derandomized fairly recently (2002) to obtain the celebrated deterministic polynomial-time algorithm for PRIMALITY TESTING. See Problem 2.5.

2.2 The Computational Model and Complexity Classes

2.2.1 Models of Randomized Computation

To develop a rigorous theory of randomized algorithms, we need to use a precise model of computation. There are several possible ways to augmenting a standard deterministic computational model (e.g. Turing machine or RAM model), such as:

- (1) The algorithm has access to a “coin-tossing black box” that provides it with (unbiased and independent) random bits on request, with each request taking one time step. This is the model we will use.
- (2) The algorithm has access to a black box that, given a number n in binary, returns a number chosen uniformly at random from $\{1, \dots, n\}$. This model is often more convenient for describing algorithms. Problem 2.2 shows that it is equivalent to Model 1, in the sense that any problem that can be solved in polynomial time on one model can also be solved in polynomial time on the other.

- (3) The algorithm is provided with an infinite tape (i.e. sequence of memory locations) that is initially filled with random bits. For polynomial-time algorithms, this is equivalent to Model 1. However, for space-bounded algorithms, this model seems stronger, as it provides the algorithm with free storage of its random bits (i.e. not counted towards its working memory). Model 1 is considered to be the “right” model for space-bounded algorithms. It is equivalent to allowing the algorithm *one-way* access to an infinite tape of random bits.

We interchangeably use “random bits” and “random coins” to refer to an algorithm’s random choices.

2.2.2 Complexity Classes

We will now define complexity classes that capture the power of efficient randomized algorithms. As is common in complexity theory, these classes are defined in terms of decision problems, where the set of inputs where the answer should be “yes” is specified by a *language* $L \subset \{0, 1\}^*$. However, the definitions generalize in natural ways to other types of computational problems, such as computing functions or solving search problems.

Recall that we say a deterministic algorithm A runs in time $t : \mathbb{N} \rightarrow \mathbb{N}$ if A takes at most $t(|x|)$ steps on every input x , and it runs in *polynomial time* if it runs time $t(n) = O(n^c)$ for a constant c . Polynomial time is a theoretical approximation to feasible computation, with the advantage that it is robust to reasonable changes in the model of computation and representation of the inputs.

Definition 2.8. \mathbf{P} is the class of languages L for which there exists a deterministic polynomial-time algorithm A such that

- $x \in L \Rightarrow A(x)$ accepts.
- $x \notin L \Rightarrow A(x)$ rejects.

For a randomized algorithm A , we say that A runs in time $t : \mathbb{N} \rightarrow \mathbb{N}$ if A takes at most $t(|x|)$ steps on every input x and every sequence of random bits.

Definition 2.9. \mathbf{RP} is the class of languages L for which there exists a probabilistic polynomial-time algorithm A such that

- $x \in L \Rightarrow \Pr[A(x) \text{ accepts}] \geq 1/2$.
- $x \notin L \Rightarrow \Pr[A(x) \text{ accepts}] = 0$.

Here (and in the definitions below) the probabilities are taken over the coin tosses of the algorithm A .

That is, \mathbf{RP} algorithms may have *false negatives*; the algorithm may sometimes say “no” even if the answer is “yes”, albeit with bounded probability. But the definition does not allow for false positives. Thus \mathbf{RP} captures efficient randomized computation with *one-sided error*. \mathbf{RP} stands

for “randomized polynomial time”. Note that the error probability of an **RP** algorithm can be reduced to $2^{-p(n)}$ for any polynomial p by running the algorithm $p(n)$ times independently and accepting the input iff at least one of the trials accepts. By the same reasoning, the $1/2$ in the definition is arbitrary, and any constant $\alpha \in (0, 1)$ or even $\alpha = 1/\text{poly}(n)$ would yield the same class of languages.

A central question in this survey is whether randomization enables us to solve more problems (e.g. decide more languages) in polynomial time:

Open Problem 2.10. Does $\mathbf{P} = \mathbf{RP}$?

Similarly, we can consider algorithms that may have false positives but no false negatives.

Definition 2.11. **co-RP** is the class of languages L whose complement \bar{L} is in **RP**. Equivalently, $L \in \mathbf{co-RP}$ if there exists a probabilistic polynomial-time algorithm A such that

- $x \in L \Rightarrow \Pr[A(x) \text{ accepts}] = 1.$
- $x \notin L \Rightarrow \Pr[A(x) \text{ accepts}] \leq 1/2.$

So, in **co-RP** we may err on NO instances, whereas in **RP** we may err on YES instances.

Using the POLYNOMIAL IDENTITY TESTING algorithm we saw earlier, we can deduce that POLYNOMIAL IDENTITY TESTING for arithmetic *formulas* is in **co-RP**. In Problem 2.4, this is generalized to arithmetic *circuits*, and thus we have:

Theorem 2.12. ARITHMETIC CIRCUIT IDENTITY TESTING over \mathbb{Z} , defined as the language

$$\text{ACIT}_{\mathbb{Z}} = \{C : C(x_1, \dots, x_n) \text{ an arithmetic circuit over } \mathbb{Z} \text{ s.t. } C = 0\},$$

is in **co-RP**.

It is common to also allow two-sided error in randomized algorithms:

Definition 2.13. **BPP** is the class of languages L for which there exists a probabilistic polynomial-time algorithm A such that

- $x \in L \Rightarrow \Pr[A(x) \text{ accepts}] \geq 2/3.$
- $x \notin L \Rightarrow \Pr[A(x) \text{ accepts}] \leq 1/3.$

Just as with **RP**, the error probability of **BPP** algorithms can be reduced from $1/3$ (or even $1/2 - 1/\text{poly}(n)$) to exponentially small by repetitions, this time taking a majority vote of the outcomes. Proving this uses some facts from probability theory, which we will review in the next section.

The cumbersome notation **BPP** stands for “bounded-error probabilistic polynomial-time,” due to the unfortunate fact that **PP** (“probabilistic polynomial-time”) refers to the definition where the

inputs in L are accepted with probability greater than $1/2$ and inputs not in L are accepted with probability at most $1/2$. Despite its name, **PP** is not a reasonable model for randomized algorithms, as it takes exponentially many repetitions to reduce the error probability. **BPP** is considered the standard complexity class associated with probabilistic polynomial-time algorithms, and thus the main question of this survey is:

Open Problem 2.14. Does $\mathbf{BPP} = \mathbf{P}$?

So far, we have considered randomized algorithms that can output an incorrect answer if they are unlucky in their coin tosses; these are called “Monte Carlo” algorithms. It is sometimes preferable to have “Las Vegas” algorithms, which always output the correct answer, but may run for a longer time if they are unlucky in their coin tosses. For this, we say that A has *expected running time* $t : \mathbb{N} \rightarrow \mathbb{N}$ if for every input x , the expectation of the number of steps taken by $A(x)$ is at most $t(|x|)$, where the expectation is taken over the coin tosses of A .

Definition 2.15. **ZPP** is the class of languages L for which there exists a probabilistic algorithm A that always decides L correctly and runs in expected polynomial time.

ZPP stands for “zero-error probabilistic polynomial time”. The following relation between **ZPP** and **RP** is left as an exercise.

Fact 2.16 (Problem 2.3). $\mathbf{ZPP} = \mathbf{RP} \cap \mathbf{co-RP}$.

We do not know any other relations between the classes associated with probabilistic polynomial time.

Open Problem 2.17. Are any of the inclusions $\mathbf{P} \subset \mathbf{ZPP} \subset \mathbf{RP} \subset \mathbf{BPP}$ proper?

One can similarly define randomized complexity classes associated with complexity measures other than time such as space or parallel computation. For example:

Definition 2.18. **RNC** is the class of languages L for which there exists a probabilistic parallel algorithm A that runs in polylogarithmic time on polynomially many processors and such that

- $x \in L \Rightarrow \Pr[A(x) \text{ accepts}] \geq 1/2$.
- $x \notin L \Rightarrow \Pr[A(x) \text{ accepts}] = 0$.

A formal model of a parallel algorithm is beyond the scope of this survey, but can be found in standard texts on algorithms or parallel computation. We have seen:

Theorem 2.19. **PERFECT MATCHING** in bipartite graphs, i.e. the language $\mathbf{PM} = \{G : G \text{ a bipartite graph with a perfect matching}\}$, is in **RNC**.

Complexity classes associated with space-bounded computation will be discussed in Section 2.4.1.

2.2.3 Tail Inequalities and Error Reduction

In the previous section, we claimed that we can reduce the error of a **BPP** algorithm by taking independent repetitions and ruling by majority vote. The intuition that this should work is based on the Law of Large Numbers: if we repeat the algorithm many times, the fraction of correct answers should approach its expectation, which is greater than $1/2$ (and thus majority rule will be correct). For complexity purposes, we need quantitative forms of this fact, which bound how many repetitions are needed to achieve a desired probability of correctness.

First, we recall a basic inequality which says that it is unlikely for (a single instantiation of) a random variable to exceed its expectation by a large factor.

Lemma 2.20 (Markov's Inequality). If X is a nonnegative real-valued random variable, then for any $\alpha > 0$,

$$\Pr[X \geq \alpha] \leq \frac{E[X]}{\alpha}$$

Markov's Inequality alone does not give very tight concentration around the expectation; to get even a 50% probability, we need to look at deviations by a factor of 2. To get tight concentration, we can take many independent copies of a random variable. There are a variety of different tail inequalities that apply for this setting; they are collectively referred to as *Chernoff Bounds*.

Theorem 2.21 (A Chernoff Bound). Let X_1, \dots, X_t be independent random variables taking values in the interval $[0, 1]$, let $X = (\sum_i X_i)/t$, and $\mu = E[X]$. Then

$$\Pr[|X - \mu| \geq \varepsilon] \leq 2 \exp(-t\varepsilon^2/4).$$

Thus, the probability that the average deviates significantly from the expectation vanishes exponentially with the number of repetitions t . We leave the proof of this Chernoff Bound as an exercise (Problem 2.7).

Now let's apply the Chernoff Bound to analyze error-reduction for **BPP** algorithms.

Proposition 2.22. The following are equivalent:

- (1) $L \in \mathbf{BPP}$.
- (2) For every polynomial p , L has a probabilistic polynomial-time algorithm with two-sided error at most $2^{-p(n)}$.
- (3) There exists a polynomial q such that L has a probabilistic polynomial-time algorithm with two-sided error at most $1/2 - 1/q(n)$.

Proof. Clearly, (2) \Rightarrow (1) \Rightarrow (3). Thus, we prove (3) \Rightarrow (2).

Given an algorithm A with error probability at most $1/2 - 1/q(n)$, consider an algorithm A' that on an input x of length n , runs A for $t(n)$ independent repetitions and rules according to the majority, where $t(n)$ is a polynomial to be determined later.

We now compute the error probability of A' on an input x of length n . Let X_i be an indicator random variable that is 1 iff the i 'th execution of $A(x)$ outputs the correct answer, and let $X = (\sum_i X_i)/t$ be the average of these indicators, where $t = t(n)$. Note that $A'(x)$ is correct when $X > 1/2$. By the error probability of A and linearity of expectations, we have $\mathbb{E}[X] \geq 1/2 + 1/q$, where $q = q(n)$. Thus, applying the Chernoff Bound with $\varepsilon = 1/q$, we have:

$$\Pr[X \leq 1/2] \leq 2 \cdot e^{-t/2q^2} < 2^{-p(n)},$$

for $t(n) = 2p(n)q(n)^2$ and sufficiently large n . □

2.3 Sampling and Approximation Problems

2.3.1 Sampling

The power of randomization is well-known to statisticians. If we want to estimate the mean of some quantity over a large population, we can do so very efficiently by taking the average over a small random sample.

Formally, here is the computational problem we are interested in solving.

Computational Problem 2.23. SAMPLING (a.k.a. $[\varepsilon]$ -APPROX ORACLE AVERAGE): Given oracle access to a function $f : \{0, 1\}^m \rightarrow [0, 1]$, estimate $\mu(f) \stackrel{\text{def}}{=} \mathbb{E}[f(U_m)]$ to within an additive error of ε , where U_m is uniformly distributed in $\{0, 1\}^m$ (as in Section 1.3). That is, output an answer in the interval $[\mu - \varepsilon, \mu + \varepsilon]$.

And here is the algorithm:

Algorithm 2.24 ($[\varepsilon]$ -APPROX ORACLE AVERAGE).

Input: oracle access to a function $f : \{0, 1\}^m \rightarrow [0, 1]$, and a desired error probability $\delta \in (0, 1)$

- (1) Choose $x_1, \dots, x_t \stackrel{\text{R}}{\leftarrow} \{0, 1\}^m$, for an appropriate choice of $t = O(\log(1/\delta)/\varepsilon^2)$.
 - (2) Query the oracle to obtain $f(x_1), \dots, f(x_t)$.
 - (3) Output $(\sum_i f(x_i))/t$.
-

The correctness of this algorithm follows from the Chernoff Bound (Theorem 2.21). Note that for constant ε and δ , the sample size t is independent of the size of the population (2^m), and we have running time $\text{poly}(m)$ even for $\varepsilon = 1/\text{poly}(m)$ and $\delta = 2^{-\text{poly}(m)}$.

For this problem, we can *prove* that no deterministic algorithm can be nearly as efficient.

Proposition 2.25. Any deterministic algorithm solving $[(1/4)]$ -APPROX ORACLE AVERAGE must make at least $2^m/2$ queries to its oracle.

Proof. Suppose we have a deterministic algorithm A that makes fewer than $2^m/2$ queries. Let Q be the set of queries made by A when all of its queries are answered by 0. (We need to specify how the queries are answered to define Q , because A make its queries adaptively, with future queries depending on how earlier queries were answered.)

Now define two functions:

$$\begin{aligned} f_0(x) &= 0 & \forall x \\ f_1(x) &= \begin{cases} 0 & x \in Q \\ 1 & x \notin Q \end{cases} \end{aligned}$$

Then A gives the same answer on both f_0 and f_1 (since all the oracle queries return 0 in both cases), but $\mu(f_0) = 0$ and $\mu(f_1) > 1/2$, so the answer must have error greater than $1/4$ for at least one of the functions. \square

Thus, randomization provides an exponential savings for approximating the average of a function on a large domain. However, this does not show that $\mathbf{BPP} \neq \mathbf{P}$. There are two reasons for this:

- (1) $[\pm\varepsilon]$ -APPROX ORACLE AVERAGE is not a decision problem, and indeed it is not clear how to define languages that capture the complexity of approximation problems. However, below we will see how a slightly more general notion of decision problem does allow us to capture approximation problems such as this one.
- (2) More fundamentally, it does not involve the standard model of input as used in the definitions of \mathbf{P} and \mathbf{BPP} . Rather than the input being a string that is explicitly given to the algorithm (where we measure complexity in terms of the length of the string), the input is an exponential-sized oracle to which the algorithm is given random access. Even though this is not the classical notion of input, it is an interesting one that has received a lot of attention in recent years, because it allows for algorithms whose running time is sublinear (or even polylogarithmic) in the actual size of the input (e.g. 2^m in the example here). As in the example here, typically such algorithms require randomization and provide approximate answers.

2.3.2 Promise Problems

Now we will try to find a variant of the $[\pm\varepsilon]$ -APPROX ORACLE AVERAGE problem that is closer to the \mathbf{P} vs. \mathbf{BPP} question. First, to obtain the standard notion of input, we consider functions that are presented in a concise form, as Boolean circuits $C : \{0, 1\}^m \rightarrow \{0, 1\}$ (analogous to the algebraic circuits defined in Section 2.1, but now the inputs take on Boolean values and the computation gates are \wedge , \vee , and \neg).

Next, we need a more general notion of decision problem than languages:

Definition 2.26. A *promise problem* Π consists of a pair (Π_Y, Π_N) of disjoint sets of strings, where Π_Y is the set of YES instances and Π_N is the set of NO instances. The corresponding computational problem is: given a string that is “promised” to be in Π_Y or Π_N , decide which is the case.

All of the complexity classes we have seen have natural promise-problem analogues, which we denote by \mathbf{prP} , \mathbf{prRP} , \mathbf{prBPP} , etc. For example:

Definition 2.27. \mathbf{prBPP} is the class of promise problems Π for which there exists a probabilistic polynomial-time algorithm A such that

- $x \in \Pi_Y \Rightarrow \Pr[A(x) \text{ accepts}] \geq 2/3$.
- $x \in \Pi_N \Rightarrow \Pr[A(x) \text{ accepts}] \leq 1/3$.

Since every language L corresponds to the promise problem (L, \bar{L}) , any fact that holds for every promise problem in some promise-class also holds for every language in the corresponding language class. In particular, if every **prBPP** algorithm can be derandomized, so can every **BPP** algorithm:

Proposition 2.28. $\text{prBPP} = \text{prP} \Rightarrow \text{BPP} = \text{P}$.

Except where otherwise noted, all of the results in this survey about language classes (like **BPP**) easily generalize to the corresponding promise classes (like **prBPP**), but we state most of the results in terms of the language classes for notational convenience.

Now we consider the following promise problem.

Computational Problem 2.29. $[\pm\varepsilon]$ -APPROX CIRCUIT AVERAGE is the promise problem CA^ε , defined as:

$$\begin{aligned} \text{CA}_Y^\varepsilon &= \{(C, p) : \mu(C) > p + \varepsilon\} \\ \text{CA}_N^\varepsilon &= \{(C, p) : \mu(C) \leq p\} \end{aligned}$$

Here ε can be a constant or a function of the input length $n = |(C, p)|$.

It turns out that this problem completely captures the power of probabilistic polynomial-time algorithms.

Theorem 2.30. For every function ε such that $1/\text{poly}(n) \leq \varepsilon(n) \leq 1 - 1/2^{n^{o(1)}}$, $[\pm\varepsilon]$ -APPROX CIRCUIT AVERAGE is **prBPP**-complete. That is, it is in **prBPP** and every promise problem in **prBPP** reduces to it.

Proof Sketch. Membership in **prBPP** follows from Algorithm 2.24 and the fact that boolean circuits can be evaluated in polynomial time.

For **prBPP**-hardness, consider any promise problem $\Pi \in \text{prBPP}$. We have a probabilistic $t(n)$ -time algorithm A that decides Π with 2-sided error at most 2^{-n} on inputs of length n , where $t(n) = \text{poly}(n)$. As an upper bound, A uses at most $t(n)$ random bits. *Thus we can view A as a deterministic algorithm on two inputs — its regular input $x \in \{0, 1\}^n$ and its coin tosses $r \in \{0, 1\}^{t(n)}$.* (This view of a randomized algorithm is useful throughout the study of pseudorandomness.) We'll write $A(x; r)$ for A 's output on input x and coin tosses r . For every n , there is a circuit $C(x; r)$ of size $\text{poly}(t(n)) = \text{poly}(n)$ that simulates the computation of A for any input $x \in \{0, 1\}^n$ $r \in \{0, 1\}^{t(n)}$, and moreover C can be constructed in time $\text{poly}(t(n)) = \text{poly}(n)$. (See any text on complexity theory for a proof that circuits can efficiently simulate Turing machine computations.) Let $C_x(r)$ be the circuit C with x hardwired in. Then the map $x \mapsto (C_x, 1/2^n)$ is a polynomial-time reduction from Π to $[\pm\varepsilon]$ -APPROX CIRCUIT AVERAGE. Indeed, if $x \in \Pi_N$, then A accepts with probability at most $1/2^n$, so $\mu(C_x) \leq 1/2^n$. And if $x \in \Pi_Y$, then $\mu(C_x) \geq 1 - 1/2^n > 1/2^n + \varepsilon(n')$, where $n' = |(C_x, 1/2^n)| = \text{poly}(n)$ and we take n sufficiently large. \square

Consequently, derandomizing this one algorithm is equivalent to derandomizing all of **prBPP**:

Corollary 2.31. $[+\varepsilon]$ -APPROX CIRCUIT AVERAGE is in **prP** if and only if **prBPP** = **prP**.

Note that our proof of Proposition 2.25, giving an exponential lower bound for $[+\varepsilon]$ -APPROX ORACLE AVERAGE does not extend to $[+\varepsilon]$ -APPROX CIRCUIT AVERAGE. Indeed, it's not even clear how to define the notion of "query" for an algorithm that is given a circuit C explicitly; it can do arbitrary computations that involve the internal structure of the circuit. Moreover, even if we restrict attention to algorithms that only use the input circuit C as if it were an oracle (other than computing the input length $|(C, p)|$ to know how long it can run), there is no guarantee that the function f_1 constructed in the proof of Proposition 2.25 has a small circuit.

2.3.3 Approximate Counting to within Relative Error

Note that $[+\varepsilon]$ -APPROX CIRCUIT AVERAGE can be viewed as the problem of approximately counting the number of satisfying assignments of a circuit $C : \{0, 1\}^m \rightarrow \{0, 1\}$ to within additive error $\varepsilon \cdot 2^m$, and a solution to this problem may give useless answers for circuits that don't have very many satisfying assignments (e.g. circuits with fewer than $2^{m/2}$ satisfying assignments). Thus people typically study approximate counting to within *relative error*. For example, given a circuit C , output a number that is within a $(1 + \varepsilon)$ factor of its number of satisfying assignments, $\#C$. Or the following essentially equivalent decision problem:

Computational Problem 2.32. $[\times(1 + \varepsilon)]$ -APPROX $\#$ CSAT is the following promise problem:

$$\begin{aligned} \text{CSAT}_Y^\varepsilon &= \{(C, N) : \#C > (1 + \varepsilon) \cdot N\} \\ \text{CSAT}_N^\varepsilon &= \{(C, N) : \#C \leq N\} \end{aligned}$$

Here ε can be a constant or a function of the input length $n = |(C, N)|$.

Unfortunately, this problem is **NP**-hard for general circuits (consider the special case that $N = 0$), so we do not expect a **prBPP** algorithm. However, there is a very pretty randomized algorithm if we restrict to DNF formulas.

Computational Problem 2.33. $[\times(1 + \varepsilon)]$ -APPROX $\#$ DNF is the restriction of $[\times(1 + \varepsilon)]$ -APPROX $\#$ CSAT to C that are formulas in *disjunctive normal form* (DNF) (i.e. an OR of clauses, where each clause is an AND of variables or their negations).

Theorem 2.34. For every function $\varepsilon(n) \geq 1/\text{poly}(n)$, $[\times(1 + \varepsilon)]$ -APPROX $\#$ DNF is in **prBPP**.

Proof. It suffices to give a probabilistic polynomial-time algorithm that estimates the number of satisfying assignments to within a $1 \pm \varepsilon$ factor. Let $\varphi(x_1, \dots, x_m)$ be the input DNF formula.

A first approach would be to apply random sampling as we have used above: Pick t random assignments uniformly from $\{0, 1\}^m$ and evaluate φ on each. If k of the assignments satisfy φ ,

output $(k/t) \cdot 2^m$. However, if $\#\varphi$ is small (e.g. $2^{m/2}$), random sampling will be unlikely to hit any satisfying assignments, and our estimate will be 0.

The idea to get around this difficulty is to embed the set of satisfying assignments, A , in a smaller set B so that sampling can be useful. Specifically, we will define sets $A' \subset B$ satisfying the following properties:

- (1) $|A'| = |A|$.
- (2) $|A'| \geq |B|/\text{poly}(n)$, where $n = |(\varphi, N)|$.
- (3) We can decide membership in A' in polynomial time.
- (4) $|B|$ computable in polynomial time.
- (5) We can sample uniformly at random from B in polynomial time.

Letting ℓ be the number of clauses, we define A' and B as follows:

$$\begin{aligned} B &= \left\{ (i, \alpha) \in [\ell] \times \{0, 1\}^m : \alpha \text{ satisfies the } i^{\text{th}} \text{ clause} \right\} \\ A' &= \left\{ (i, \alpha) \in B : \alpha \text{ does not satisfy any clauses before the } i^{\text{th}} \text{ clause} \right\} \end{aligned}$$

Now we verify the desired properties:

- (1) $|A| = |A'|$ because for each satisfying assignment α , A' contains only one pair (i, α) , namely the one where the first clause satisfied by α is the i^{th} one.
- (2) The size of A' and B can differ by at most a factor of ℓ , since each satisfying assignment α occurs at most ℓ times in B .
- (3) It is easy to decide membership in A' in polynomial time.
- (4) $|B| = \sum_{i=1}^{\ell} 2^{m-m_i}$, where m_i is the number of literals in clause i (after removing contradictory clauses that contain both a variable and its negation).
- (5) We can randomly sample from B as follows. First pick a clause with probability proportional to the number of satisfying assignments it has (2^{m-m_i}). Then, fixing those variables in the clause (e.g. if x_j is in the clause, set $x_j = 1$, and if $\neg x_j$ is in the clause, set $x_j = 0$), assign the rest of the variables uniformly at random.

Putting this together, we deduce the following algorithm:

Algorithm 2.35 ($[\times(1 + \varepsilon)]$ -APPROX $\#\text{DNF}$).

Input: a DNF formula $\varphi(x_1, \dots, x_m)$ with ℓ clauses

- (1) Generate a random sample of t points in $B = \{(i, \alpha) \in [\ell] \times \{0, 1\}^m : \alpha \text{ satisfies the } i^{\text{th}} \text{ clause}\}$, for an appropriate choice of $t = O(1/(\varepsilon/\ell)^2)$.
 - (2) Let $\hat{\mu}$ be the fraction of sample points that land in $A' = \{(i, \alpha) \in B : \alpha \text{ does not satisfy any clauses before the } i^{\text{th}} \text{ clause}\}$.
 - (3) Output $\hat{\mu} \cdot |B|$.
-

By the Chernoff bound, for an appropriate choice of $t = O(1/(\varepsilon/\ell)^2)$, we have $\hat{\mu} = |A'|/|B| \pm \varepsilon/\ell$ with high probability (where we write $\gamma = \alpha \pm \beta$ to mean that $\gamma \in [\alpha - \beta, \alpha + \beta]$). Thus, with high probability the output of the algorithm satisfies:

$$\begin{aligned}\hat{\mu} \cdot |B| &= |A'| \pm \varepsilon|B|/\ell \\ &= |A| \pm \varepsilon|A|.\end{aligned}$$

□

There is no deterministic polynomial-time algorithm known for this problem:

Open Problem 2.36. Give a deterministic polynomial-time algorithm for approximately counting the number of satisfying assignments to a DNF formula.

However, when we study pseudorandom generators in Chapter 7, we will see a quasipolynomial-time derandomization of the above algorithm (i.e. one in time $2^{\text{polylog}(n)}$).

2.3.4 MAXCUT

We give an example of another algorithmic problem for which random sampling is a useful tool.

Definition 2.37. For a graph $G = (V, E)$ and $S, T \subset V$, define $\text{cut}(S, T) = \{\{u, v\} \in E : u \in S, v \in T\}$, and $\text{cut}(S) = \text{cut}(S, V \setminus S)$.

Computational Problem 2.38. MAXCUT (search version): Given a graph G , find a largest cut in G , i.e. a set S maximizing $|\text{cut}(S)|$.

Solving this problem optimally is **NP**-hard (in contrast to MINCUT, which is known to be in **P**). However, there is a simple randomized algorithm that finds a cut of expected size at least $|E|/2$ (which is of course at least $1/2$ the optimal, and hence this is referred to as a “ $1/2$ -approximation algorithm”):

Algorithm 2.39 (MAXCUT approximation).

Input: a graph $G = (V, E)$ (with no self-loops)

Output a random subset $S \subset V$. That is, place each vertex v in S independently with probability $1/2$.

To analyze this algorithm, consider any edge $e = (u, v)$. Then the probability that e crosses the cut is $1/2$. By linearity of expectations, we have:

$$\mathbb{E}[|\text{cut}(S)|] = \sum_{e \in E} \Pr[e \text{ is cut}] = |E|/2.$$

This also serves as a proof, via the probabilistic method, that every graph (without self-loops) has a cut of size at least $|E|/2$.

In Chapter 3, we will see how to derandomize this algorithm. We note that there is a much more sophisticated randomized algorithm that finds a cut whose expected size is within a factor of .878 of the largest cut in the graph (and this algorithm can also be derandomized).

2.4 Random Walks and S-T CONNECTIVITY

2.4.1 Graph Connectivity

One of the most basic problems in computer science is that of deciding connectivity in graphs:

Computational Problem 2.40. S-T CONNECTIVITY: Given a directed graph G and two vertices s and t , is there a path from s to t in G ?

This problem can of course be solved in linear time using breadth-first or depth-first search. However, these algorithms also require linear space. It turns out that S-T CONNECTIVITY can in fact be solved using much less workspace. (When measuring the space complexity of algorithms, we do not count the space for the (read-only) input and (write-only) output.)

Theorem 2.41. There is a deterministic algorithm deciding S-T CONNECTIVITY using space $O(\log^2 n)$ (and time $O(n^{\log n})$).

Proof. The following recursive algorithm $\text{IsPath}(G, u, v, k)$ decides whether there is a path of length at most k from u to v .

Algorithm 2.42 (Recursive S-T CONNECTIVITY).

$\text{IsPath}(G, u, v, k)$:

- (1) If $k = 0$, accept if $u = v$.
- (2) If $k = 1$, accept if $u = v$ or (u, v) is an edge in G .
- (3) Otherwise, loop through all vertices w in G and accept if both $\text{IsPath}(G, u, w, \lceil k/2 \rceil)$ and $\text{IsPath}(G, w, v, \lfloor k/2 \rfloor)$ accept for some w .

We can solve S-T CONNECTIVITY by running $\text{IsPath}(G, s, t, n)$, where n is the number of vertices in the graph. The algorithm has $\log n$ levels of recursion and uses $\log n$ space per level of recursion (to store the vertex w), for a total space bound of $\log^2 n$. Similarly, the algorithm uses linear time per level of recursion, for a total time bound of $O(n)^{\log n}$. \square

It is not known how to improve the space bound in Theorem 2.41 or to get the running time down to polynomial while maintaining space $n^{o(1)}$. For *undirected graphs*, however, we can do much better using a randomized algorithm. Specifically, we can place the problem in the following class:

Definition 2.43. A language L is in **RL** if there exists a randomized algorithm A that always halts, uses space at most $O(\log n)$ on inputs of length n , and satisfies:

- $x \in L \Rightarrow \Pr[A(x) \text{ accepts}] \geq 1/2$.
- $x \notin L \Rightarrow \Pr[A(x) \text{ accepts}] = 0$.

Recall that our model of a randomized space-bounded machine is one that has access to a coin-tossing box (rather than an infinite tape of random bits), and thus must explicitly store in its workspace any random bits it needs to remember. The requirement that A always halts ensures that its running time is at most $2^{O(\log n)} = \text{poly}(n)$, because otherwise there would be a loop in its configuration space. Similarly to **RL**, we can define **L** (deterministic logspace), **co-RL** (one-sided error with errors only on NO instances), and **BPL** (two-sided error).

Now we can state the theorem for undirected graphs.

Computational Problem 2.44. **UNDIRECTED S-T CONNECTIVITY:** Given an undirected graph G and two vertices s and t , is there a path from s to t in G ?

Theorem 2.45. **UNDIRECTED S-T CONNECTIVITY** is in **RL**.

Proof Sketch. The algorithm simply does a polynomial-length random walk starting at s .

Algorithm 2.46 (UNDIRECTED S-T CONNECTIVITY via Random Walks).

Input: (G, s, t) , where $G = (V, E)$ has n vertices.

- (1) Let $v = s$.
- (2) Repeat $\text{poly}(n)$ times:
 - (a) If $v = t$, halt and accept.
 - (b) Else let $v \stackrel{R}{\leftarrow} \{w : (v, w) \in E\}$.
- (3) Reject (if we haven't visited t yet).

Notice that this algorithm only requires space $O(\log n)$, to maintain the current vertex v as well as a counter for the number of steps taken. Clearly, it never accepts when there isn't a path from s to t . In the next section, we will prove that in any connected undirected graph, a random walk of length $\text{poly}(n)$ from one vertex will hit any other vertex with high probability. Applying this to the connected component containing s , it follows that the algorithm accepts with high probability when s and t are connected. \square

This algorithm, dating from the 1970's, was derandomized only in 2005. We will cover the derandomized algorithm in Section 4.4. However, the general question of derandomizing space-bounded algorithms remains open.

Open Problem 2.47. Does $\text{RL} = \text{L}$? Does $\text{BPL} = \text{L}$?

2.4.2 Random Walks on Graphs

For generality that will be useful later, many of the definitions in this section will be given for *directed multigraphs* (which we will refer to as *digraphs* for short). By multigraph, we mean that we allow G to have parallel edges and self-loops. Henceforth, we will refer to graphs without parallel edges and self-loops as *simple graphs*. We call a digraph d -regular if every vertex has indegree d and outdegree d . A self-loop is considered a single directed edge (i, i) from a vertex i to itself, so contributes 1 to both the indegree and outdegree of vertex i . An undirected graph is a graph where the number of edges from i to j equals the number of edges from j to i for every i and j . (When $i \neq j$, we think of a pair of edges (i, j) and (j, i) as comprising a single undirected edge $\{i, j\}$, and a self-loop (i, i) also corresponds to the single undirected edge $\{i\}$.)

To analyze the random-walk algorithm of the previous section, it suffices to prove a bound on the *hitting time* of random walks.

Definition 2.48. For a digraph $G = (V, E)$, we define its *hitting time* as

$$\text{hit}(G) = \max_{i, j \in V} \min\{t : \Pr[\text{a random walk of length } t \text{ started at } i \text{ visits } j] \geq 1/2\}.$$

We note that $\text{hit}(G)$ is often defined as the maximum over vertices i and j of the *expected* time for a random walk from i to visit j . The two definitions are the same up to a factor of 2, and the above is more convenient for our purposes.

We will prove:

Theorem 2.49. For every connected undirected graph G with n vertices and maximum degree d , we have $\text{hit}(G) = O(d^2 n^3 \log n)$.

We observe that it suffices to prove this theorem for d -regular graphs, because any undirected graph can be made regular by adding self-loops (and this can only increase the hitting time). This is the part of our proof that fails for directed graphs (adding self-loops cannot correct an imbalance between indegree and outdegree at a vertex), and indeed Problem 2.10 shows that general directed graphs can have exponential hitting time.

There are combinatorial methods for proving the above theorem, but we will prove it using a linear-algebraic approach, as the same methods will be very useful in our study of expander graphs. For an n -vertex digraph G , we define its *random-walk transition matrix*, or *random-walk matrix* for short, to be the $n \times n$ matrix M where $M_{i,j}$ is the probability of going from vertex i to vertex j in one step. That is, $M_{i,j}$ is the number of edges from i to j divided by the outdegree of i . In case G is d -regular, M is simply the adjacency matrix of G divided by d . Notice that for every probability distribution $\pi \in \mathbb{R}^n$ on the vertices of G (written as a row vector), the vector πM is the probability distribution obtained by selecting a vertex i according to π and then taking one step of the random walk to end at a vertex j . This is because $(\pi M)_j = \sum_i \pi_i M_{i,j}$.

In our application to bounding $\text{hit}(G)$ for a *regular* digraph G , we start at a probability distribution $\pi = (0, \dots, 0, 1, 0, \dots, 0) \in \mathbb{R}^n$ concentrated at vertex i , and are interested in the distribution πM^t we get after taking t steps on the graph. Specifically, we'd like to show that it places nonnegligible mass on vertex j for $t = \text{poly}(n)$. We will do this by showing that it in fact converges to the uniform distribution $u = (1/n, 1/n, \dots, 1/n) \in \mathbb{R}^n$ within a polynomial number of steps. Note that $uM = u$ by the regularity of G , so convergence to u is possible (and will be guaranteed given some additional conditions on G).

We will measure the rate of convergence in the ℓ_2 norm. For vectors $x, y \in \mathbb{R}^n$, we will use the standard inner product $\langle x, y \rangle = \sum_i x_i y_i$, and ℓ_2 norm $\|x\| = \sqrt{\langle x, x \rangle}$. We write $x \perp y$ to mean that x and y are orthogonal, i.e. $\langle x, y \rangle = 0$. We want to determine how large k needs to be so that $\|\pi M^k - u\|$ is “small”. This is referred to as the *mixing time* of the random walk. Mixing time can be defined with respect to various distance measures and the ℓ_2 norm is not the most natural one, but it has the advantage that we will be able to show that the distance decreases noticeably in every step. This is captured by the following quantity.

Definition 2.50. For a regular digraph G with random-walk matrix M , we define

$$\lambda(G) \stackrel{\text{def}}{=} \max_{\pi} \frac{\|\pi M - u\|}{\|\pi - u\|} = \max_{x \perp u} \frac{\|xM\|}{\|x\|},$$

where the first maximization is over all *probability distributions* $\pi \in [0, 1]^n$ and the second is over all vectors $x \in \mathbb{R}^n$ such that $x \perp u$. We write $\gamma(G) \stackrel{\text{def}}{=} 1 - \lambda(G)$.

To see that the first definition of $\lambda(G)$ is smaller than or equal to the second, note that for any probability distribution π , the vector $x = (\pi - u)$ is orthogonal to uniform (i.e. the sum of its entries is zero). For the converse, observe that given any vector $x \perp u$, the vector $\pi = u + \alpha x$ is a probability distribution for a sufficiently small α . It can be shown that $\lambda(G) \in [0, 1]$. (This follows from Problems 2.10 and 2.9.)

The following lemma is immediate from the definition of $\lambda(G)$.

Lemma 2.51. Let G be a regular digraph with random-walk matrix M . For every initial probability distribution π on the vertices of G and every $t \in \mathbb{N}$, we have

$$\|\pi M^t - u\| \leq \lambda(G)^t \cdot \|\pi - u\| \leq \lambda(G)^t.$$

Proof. The first inequality follows from the definition of $\lambda(G)$ and induction. For the second, we have:

$$\begin{aligned} \|\pi - u\|^2 &= \|\pi\|^2 + \|u\|^2 - 2\langle \pi, u \rangle \\ &= \sum_i \pi_i^2 + 1/N - 2/N \sum_i \pi_i \\ &\leq \sum_i \pi_i + 1/N - 2/N \sum_i \pi_i \\ &= 1 - 1/N \leq 1. \end{aligned}$$

□

Thus a smaller value of $\lambda(G)$ (equivalently, a larger value of $\gamma(G)$) means that the random walk mixes more quickly. Specifically, for $t \geq \ln(n/\varepsilon)/\gamma(G)$, it follows that every entry of πM^t has probability mass at least $1/n - (1 - \gamma(G))^t \geq (1 - \varepsilon)/n$, and thus we should think of the walk as being “mixed” within $O((\log n)/\gamma(G))$. Note that after $O(1/\gamma(G))$ steps, the ℓ_2 distance is already small (say at most $1/10$), but this is not a satisfactory notion of mixing — a distribution that assigns ε^2 mass to $1/\varepsilon^2$ vertices has ℓ_2 distance smaller than ε from the uniform distribution.

Corollary 2.52. For every regular digraph G , $\text{hit}(G) = O(n \log n / \gamma(G))$.

Proof. Let i, j be any two vertices of G . As argued above, a walk of length $\ln(2n)/\gamma(G)$ from any start vertex has a probability of at least $1/2n$ of ending at j . Thus, if we do $O(n)$ such walks consecutively, we will hit j with probability at least $1/2$. \square

Thus, our task reduces to bounding $\gamma(G)$:

Theorem 2.53. If G is a connected, nonbipartite, and d -regular undirected graph on n vertices, then $\gamma(G) = \Omega(1/(dn)^2)$.

Theorem 2.53 and Corollary 2.52 imply Theorem 2.49, as any undirected and connected graph of maximum degree d can be made $(d+1)$ -regular and nonbipartite by adding self-loops to each vertex. (We remark that the bounds of Theorems 2.53 and 2.49 are not tight.) Theorem 2.53 is proven in Problem 2.9, using a connection with eigenvalues described in the next section.

2.4.3 Eigenvalues

Recall that a nonzero vector $v \in \mathbb{R}^n$ is an *eigenvector* of $n \times n$ matrix M if $vM = \lambda v$ for some $\lambda \in \mathbb{R}$, which is called the corresponding *eigenvalue*. A useful feature of *symmetric* matrices is that they can be described entirely in terms of their eigenvectors and eigenvalues.

Theorem 2.54 (Spectral Theorem for Symmetric Matrices). If M is a symmetric $n \times n$ real matrix with distinct eigenvalues μ_1, \dots, μ_k , then the eigenspaces $W_i = \{v \in \mathbb{R}^n : vM = \mu_i M\}$ are orthogonal (i.e. $v \in W_i, w \in W_j \Rightarrow v \perp w$ if $i \neq j$) and span \mathbb{R}^n (i.e. $\mathbb{R}^n = W_1 + \dots + W_k$). We refer to the dimension of W_i as the *multiplicity* of eigenvalue μ_i . In particular, \mathbb{R}^n has a basis consisting of orthogonal eigenvectors v_1, \dots, v_n having respective eigenvalues $\lambda_1, \dots, \lambda_n$, where the number of times μ_i occurs among the λ_j 's exactly equals the multiplicity of μ_i .

Notice that if G is a undirected regular graph, then its random-walk matrix M is symmetric. We know that $uM = u$, so the uniform distribution is an eigenvector of eigenvalue 1. Let v_2, \dots, v_n and $\lambda_2, \dots, \lambda_n$ be the remaining eigenvectors and eigenvalues, respectively. Given any probability distribution π , we can write it as $\pi = u + c_2 v_2 + \dots + c_n v_n$. Then the probability distribution after k steps on the random walk is

$$\pi M^t = u + \lambda_2^t c_2 v_2 + \dots + \lambda_n^t c_n v_n.$$

In Problem 2.9, it is shown that all of the λ_i 's have absolute value at most 1. Notice that if they all have magnitude strictly smaller than 1, then πM^t indeed converges to u . Thus it is not surprising that our measure of mixing rate, $\lambda(G)$, equals the absolute value of the second largest eigenvalue (in magnitude).

Lemma 2.55. Let G be an undirected graph with random-walk matrix M . Let $\lambda_1, \dots, \lambda_n$ be the eigenvalues of M , sorted so that $1 = \lambda_1 \geq |\lambda_2| \geq |\lambda_3| \geq \dots \geq |\lambda_n|$. Then $\lambda(G) = |\lambda_2|$.

Proof. Let $u = v_1, v_2, \dots, v_n$ be the basis of orthogonal eigenvectors corresponding to the λ_i 's. Given any vector $x \perp u$, we can write $x = c_2 v_2 + \dots + c_n v_n$. Then:

$$\begin{aligned} \|xM\|^2 &= \|\lambda_2 c_2 v_2 + \dots + \lambda_n c_n v_n\|^2 \\ &= \lambda_2^2 c_2^2 \|v_2\|^2 + \dots + \lambda_n^2 c_n^2 \|v_n\|^2 \\ &\leq |\lambda_2|^2 \cdot (c_2^2 \|v_2\|^2 + \dots + c_n^2 \|v_n\|^2) \\ &= |\lambda_2|^2 \cdot \|x\|^2 \end{aligned}$$

Equality is achieved with $x = v_2$. □

Thus, bounding $\lambda(G)$ amounts to bounding the eigenvalues of G . Due to this connection, $\gamma(G) = 1 - \lambda(G)$ is often referred to as the *spectral gap*, as it is the gap between the largest eigenvalue and the second largest eigenvalue in absolute value.

2.4.4 Markov Chain Monte Carlo

Random walks are a powerful tool in the design of randomized algorithms. In particular, they are the heart of the “Markov Chain Monte Carlo” method, which is widely used in statistical physics and for solving approximate counting problems. In these applications, the goal is to generate a random sample from an *exponentially* large space, such as an (almost) uniformly random perfect matching for a given bipartite graph G . (It turns out that this is equivalent to approximately counting the number of perfect matchings in G .) The approach is to do a random walk on an appropriate (regular) graph \hat{G} defined on the space (e.g. by doing random local changes on the current perfect matching). Even though \hat{G} is typically of size exponential in the input size $n = |G|$, in many cases it can be proven to have mixing time $\text{poly}(n) = \text{polylog}(|\hat{G}|)$, a property referred to as *rapid mixing*. These Markov Chain Monte Carlo methods provide some of the best examples of problems where randomization yields algorithms that are exponentially faster than all known deterministic algorithms.

2.5 Exercises

Problem 2.1 (Schwartz–Zippel lemma). Prove Lemma 2.4: If $f(x_1, \dots, x_n)$ is a nonzero polynomial of degree d over a field (or integral domain) \mathbb{F} and $S \subset \mathbb{F}$, then

$$\Pr_{\alpha_1, \dots, \alpha_n \xleftarrow{\mathbb{R}} S} [f(\alpha_1, \dots, \alpha_n) = 0] \leq \frac{d}{|S|}.$$

You may use the fact that every nonzero *univariate* polynomial of degree d over \mathbb{F} has at most d roots.

Problem 2.2 (Robustness of the model). Suppose we modify our model of randomized computation to allow the algorithm to obtain a random element of $\{1, \dots, m\}$ for any number m whose binary representation it has already computed (as opposed to just allowing it access to random bits). Show that this would not change the classes **BPP** and **RP**.

Problem 2.3 (Zero error vs. 1-sided error). Prove that **ZPP** = **RP** \cap **co-RP**.

Problem 2.4 (POLYNOMIAL IDENTITY TESTING for integer circuits). In this problem, you will show how to do POLYNOMIAL IDENTITY TESTING for arithmetic *circuits* over the integers. The Prime Number Theorem says that the number of primes less than T is $(1 \pm o(1)) \cdot T / \ln T$, where the $o(1)$ tends to 0 as $T \rightarrow \infty$. You may use this fact in the problem below.

- (1) Show that if N is a nonzero integer and $M \stackrel{R}{\leftarrow} \{1, \dots, \log^2 N\}$, then

$$\Pr[N \not\equiv 0 \pmod{M}] = \Omega(1/\log \log N).$$

- (2) Use the above to prove Theorem 2.12: ARITHMETIC CIRCUIT IDENTITY TESTING over \mathbb{Z} is in **co-RP**.
-

Problem 2.5 (POLYNOMIAL IDENTITY TESTING via Modular Reduction). In this problem, you will analyze an alternative to the algorithm seen in class, which directly handles polynomials of degree larger than the field size. It is based on the same idea as Problem 2.4, using the fact that polynomials over a field have many of the same algebraic properties as the integers.

The following definitions and facts may be useful: A polynomial $f(x)$ over a field \mathbb{F} is called *irreducible* if it has no nontrivial factors (i.e. factors other than constants from \mathbb{F} or constant multiples of f). Analogously to prime factorization of integers, every polynomial over \mathbb{F} can be factored into irreducible polynomials and this factorization is unique (up to reordering and constant multiples). It is known that the number of irreducible polynomials of degree at most d over a field \mathbb{F} is at least $|\mathbb{F}|^{d+1}/2d$. (This is similar to the Prime Number Theorem for integers mentioned in Problem 2.4, but is easier to prove.) For polynomials $f(x)$ and $g(x)$, $f(x) \bmod g(x)$ is the remainder when f is divided by g . (More background on polynomials over finite fields can be found in the references listed in Section 2.6.)

In this problem, we consider a version of the POLYNOMIAL IDENTITY TESTING problem where a polynomial $f(x_1, \dots, x_n)$ over finite field \mathbb{F} is presented as a formula built up from elements of \mathbb{F} and the variables x_1, \dots, x_n using addition, multiplication, and *exponentiation* with exponents given in *binary*. We also assume that we are given a representation of \mathbb{F} enabling addition, multiplication, and division in \mathbb{F} to be done quickly.

- (1) Let $f(x)$ be a univariate polynomial of degree $\leq D$ over a field \mathbb{F} . Prove that there are constants c, c' such that if $f(x)$ is nonzero (as a formal polynomial) and $g(x)$ is a randomly selected polynomial of degree at most $d = c \log D$, then the probability that $f(x) \bmod g(x)$ is nonzero is at least $1/c' \log D$. Deduce a randomized, polynomial-time identity test for *univariate* polynomials presented in the above form.
 - (2) Obtain a randomized polynomial-time identity test for multivariate polynomials by giving a (deterministic) reduction to the univariate case.
-

Problem 2.6 (PRIMALITY TESTING). (1) Show that for every positive integer n , the polynomial identity $(x + 1)^n \equiv x^n + 1 \pmod{n}$ holds iff n is prime.

(2) Obtain a **co-RP** algorithm for the language $\text{PRIMALITY TESTING} = \{n : n \text{ prime}\}$ using Part 1 together with Problem 2.5. (In your analysis, remember that the integers modulo n are a field only when n is prime.)

Problem 2.7 (Chernoff Bound). Let X_1, \dots, X_t be independent $[0, 1]$ -valued random variables, and $X = \sum_{i=1}^t X_i$. (Note that, in contrast to the statement of Theorem 2.21, here we are writing X for the sum of the X_i 's rather than their average.)

- (1) Show that for every $r \in [0, 1/2]$, $\mathbb{E}[e^{rX}] \leq e^{r\mathbb{E}[X] + r^2 t}$. (Hint: $1 + x \leq e^x \leq 1 + x + x^2$ for all $x \in [0, 1/2]$.)
 - (2) Deduce the Chernoff Bound of Theorem 2.21: $\Pr[X \geq \mathbb{E}[X] + \varepsilon t] \leq e^{-\varepsilon^2 t/4}$ and $\Pr[X \leq \mathbb{E}[X] - \varepsilon t] \leq e^{-\varepsilon^2 t/4}$.
 - (3) Where did you use the independence of the X_i 's?
-

Problem 2.8 (Necessity of Randomness for Identity Testing*). In this problem, we consider the “oracle version” of the identity testing problem, where an arbitrary polynomial $f : \mathbb{F}^m \rightarrow \mathbb{F}$ of degree d is given as an oracle and the problem is to test whether $f = 0$. Show that any deterministic algorithm that solves this problem when $m = d = n$ must make at least 2^n queries to the oracle (in contrast to the randomized identity testing algorithm from class, which makes only one query provided that $|\mathbb{F}| \geq 2n$).

Is this a proof that $\mathbf{P} \neq \mathbf{RP}$? Explain.

Problem 2.9 (Spectral Graph Theory). Let M be the random-walk matrix for a d -regular *undirected* graph $G = (V, E)$ on n vertices. We allow G to have self-loops and multiple edges. Recall that the uniform distribution is an eigenvector of M of eigenvalue $\lambda_1 = 1$. Prove the following statements. (Hint: for intuition, it may help to think about what the statements mean for the behavior of the random walk on G .)

- (1) All eigenvalues of M have absolute value at most 1.
- (2) G is disconnected \iff 1 is an eigenvalue of multiplicity at least 2.
- (3) Suppose G is connected. Then G is bipartite \iff -1 is an eigenvalue of M .
- (4) G connected \implies all eigenvalues of M other than λ_1 are at most $1 - 1/\text{poly}(n, d)$. To do this, it may help to first show that the second largest eigenvalue of M (not necessarily in absolute value) equals

$$\max_x \langle xM, x \rangle = 1 - \frac{1}{d} \cdot \min_x \sum_{(i,j) \in E} (x_i - x_j)^2,$$

where the maximum/minimum is taken over all vectors x of length 1 such that $\sum_i x_i = 0$, and $\langle x, y \rangle = \sum_i x_i y_i$ is the standard inner product. For intuition, consider restricting the above maximum/minimum to $x \in \{+\alpha, -\beta\}^n$ for $\alpha, \beta > 0$.

- (5) G connected and nonbipartite \implies all eigenvalues of M (other than 1) have absolute value at most $1 - 1/\text{poly}(n, d)$ and thus $\gamma(G) \geq 1/\text{poly}(n, d)$.
- (6*) Establish the (tight) bound $1 - \Omega(1/d \cdot D \cdot n)$ in Part 4, where D is the diameter of the graph. Conclude that $\gamma(G) = \Omega(1/d^2 n^2)$ if G is connected and nonbipartite.

Problem 2.10 (Hitting Time and Eigenvalues for Directed Graphs).

- (1) Show that for every n , there exists a digraph G with n vertices, outdegree 2, and $\text{hit}(G) = 2^{\Omega(n)}$.
- (2) Let G be a *regular* digraph with random-walk matrix M . Show that $\lambda(G) = \sqrt{\lambda(G')}$, where G' is the *undirected* graph whose random-walk matrix is MM^T .
- (3) A digraph G is called *Eulerian* if it is connected and every vertex has the same indegree as outdegree.³ (For example, if we take an connected undirected graph and replace each undirected edge $\{u, v\}$ with the two directed edges (u, v) and (v, u) , we obtain an Eulerian digraph.) Show that if G is an n -vertex Eulerian digraph of maximum degree d , then $\text{hit}(G) = \text{poly}(n, d)$.

Problem 2.11 (Consequences of Derandomizing prBPP). Even though **prBPP** is a class of decision problems, it also captures many other types of problems that can be solved by randomized algorithms:

- (1) (Approximation Problems) Consider the task of approximating a function $f : \{0, 1\}^* \rightarrow \mathbb{N}$ to within a $(1 + \varepsilon(n))$ multiplicative factor, for some function $\varepsilon : \mathbb{N} \rightarrow \mathbb{R}^{\geq 0}$. Suppose there is a probabilistic polynomial-time algorithm A such that for all x , $\Pr[A(x) \leq f(x) \leq$

³This terminology comes from the fact that these are precisely the digraphs that have *Eulerian circuits*, closed paths that visit all vertices use every directed edge exactly once.

$(1 + \varepsilon(|x|)) \cdot A(x)] \geq 2/3$, where the probability is taken over the coin tosses of A . Show that if $\mathbf{prBPP} = \mathbf{prP}$, then there is a deterministic polynomial-time algorithm B such that for all x , $B(x) \leq f(x) \leq (1 + \varepsilon(|x|)) \cdot B(x)$. (Hint: use a definition similar to Computational Problem 2.32.)

- (2) (**NP** Search Problems) An **NP** search problem is specified by a polynomial-time verifier V and a polynomial p ; the problem is, given an input $x \in \{0, 1\}^n$, find a string $y \in \{0, 1\}^{p(n)}$ such that $V(x, y) = 1$. Suppose that such a search problem can be solved in probabilistic polynomial time, i.e. there is a probabilistic polynomial-time algorithm A such that for every input $x \in \{0, 1\}^n$, outputs $y \in \{0, 1\}^{p(n)}$ such that $V(x, y) = 1$ with probability at least $2/3$ over the coin tosses of A . Show that if $\mathbf{prBPP} = \mathbf{prP}$, then there is a deterministic polynomial-time algorithm B such that for every $x \in \{0, 1\}^n$, $B(x)$ always outputs $y \in \{0, 1\}^{p(n)}$ such that $V(x, y) = 1$. (Hint: consider a promise problem whose instances are pairs (x, r) where r is a prefix of the coin tosses of A .)
- (3) (**MA** Search Problems) In a relaxation of **NP**, known as **MA**, we allow the verifier V to be probabilistic. Let $\Pi = (\Pi_Y, \Pi_N)$ be a promise problem decided by V , i.e. $\Pr[V(x, y) = 1] \geq 2/3$ for $(x, y) \in \Pi_Y$, and $\Pr[V(x, y) = 1] \leq 1/3$ for $(x, y) \in \Pi_N$. Suppose that the corresponding search problem can be solved in probabilistic polynomial time, i.e. there is a probabilistic polynomial-time algorithm A such that for every input $x \in \{0, 1\}^n$, outputs $y \in \{0, 1\}^{p(n)}$ such that $(x, y) \in \Pi_Y$ with probability at least $2/3$ over the coin tosses of A . Show that if $\mathbf{prBPP} = \mathbf{prP}$, then there is a deterministic polynomial-time algorithm B such that for every $x \in \{0, 1\}^n$, $B(x)$ always outputs $y \in \{0, 1\}^{p(n)}$ such that $(x, y) \notin \Pi_N$. (Note that this conclusion is equivalent to $(x, y) \in \Pi_Y$ if Π is a language, but is weaker otherwise.)
- (4) Using Part 3, the Prime Number Theorem (see Problem 2.4), and the fact that **PRIMALITY TESTING** is in **BPP** (Problem 2.6) to show that if $\mathbf{prBPP} = \mathbf{prP}$, then there is a deterministic polynomial-time algorithm that given a number N , outputs a prime in the interval $[N, 2N]$ for all sufficiently large N .

2.6 Chapter Notes and References

Recommended textbooks on randomized algorithms are Motwani–Raghavan [MR] and Mitzenmacher–Upfal [MU]. The algorithmic power of randomization became apparent in the 1970’s with a number of striking examples, notably Berlekamp’s algorithm for factoring polynomials over large finite fields [Ber3] and the Miller–Rabin [Mil1, Rab] and Solovay–Strassen [SS] algorithms for **PRIMALITY TESTING**. The randomized algorithm for **POLYNOMIAL IDENTITY TESTING** was independently discovered by DeMillo and Lipton [DL], Schwartz [Sch], and Zippel [Zip]. More recently, a deterministic polynomial-time **POLYNOMIAL IDENTITY TESTING** algorithm for formulas in $\Sigma\Pi\Sigma$ form with a constant number of terms was given by Kayal and Saxena [KS], improving a previous quasipolynomial-time algorithm of Dvir and Shpilka [DS]. Problem 2.4 is from [Sch]. Problem 2.8 is from [LV1]. Algorithms for **POLYNOMIAL IDENTITY TESTING** and the study of arithmetic circuits more generally are covered in the recent survey by Shpilka and Yehudayoff [SY]. Recommended textbooks on abstract algebra and finite fields are [Art, LN].

The randomized algorithm for **PERFECT MATCHING** is due to Lovász [Lov1], who also showed

how to extend the algorithm to nonbipartite graphs. An efficient parallel randomized algorithm for *finding* a perfect matching was given by Karp, Upfal, and Wigderson [KUW] (see also [MVV]). A randomized algorithm for finding a perfect matching in the same sequential time complexity as Lovász’s algorithm was given recently by Mucha and Sankowski [MS] (see also [Har]).

The efficient sequential and parallel algorithms for DETERMINANT mentioned in the text are due [Str, CW2] and [Csa, BvzGH, Ber1], respectively. For more on algebraic complexity and parallel algorithms, we refer to the textbooks [BCS] and [Lei], respectively. The POLYNOMIAL IDENTITY TESTING and PRIMALITY TESTING algorithms of Problems 2.5 and 2.6 are due to Agrawal and Biswas [AB]. Agrawal, Kayal, and Saxena [AKS] derandomized the PRIMALITY TESTING algorithm of [AB] to prove that PRIMALITY TESTING is in \mathbf{P} .

The randomized complexity classes \mathbf{RP} , \mathbf{BPP} , \mathbf{ZPP} , and \mathbf{PP} were formally defined by Gill [Gil2], who conjectured that $\mathbf{BPP} \neq \mathbf{P}$ (in fact $\mathbf{ZPP} \neq \mathbf{P}$). Chernoff Bounds are named after H. Chernoff [Che2]; the version in Theorem 2.21 is due to Hoeffding [Hoe] and is sometimes referred to as Hoeffding’s Inequality. The monograph by Dubhashi and Panconesi [DP] has a detailed coverage of Chernoff Bounds and other tail inequalities. Problem 2.3 is due to Rabin (cf. [Gil2]).

The computational perspective on sampling, as introduced in Section 2.3.1, is surveyed in [Gol1, Gol2]. SAMPLING is perhaps the simplest example of a computational problem where randomization enables algorithms with running time sublinear in the size of the input. Such *sublinear-time algorithms* are now known for a wide variety of interesting computational problems; see the surveys [Ron, Rub].

Promise problems were introduced by Even, Selman, and Yacobi [ESY]. For survey of their role in complexity theory, see Goldreich [Gol5]. Problem 2.11 on the consequences of $\mathbf{prBPP} = \mathbf{prP}$ is from [Gol9].

The randomized algorithm for $[\times(1 + \varepsilon)]$ -APPROX #DNF is due to Karp and Luby [KLM], who initiated the study of randomized algorithms for approximate counting problems. A $1/2$ -approximation algorithm for MAXCUT was first given in [SG]; that algorithm can be viewed as a natural derandomization of Algorithm 2.39. (See Algorithm 3.17.) The .878-approximation algorithm was given by Goemans and Williamson [GW].

The $O(\log^2 n)$ -space algorithm for S-T CONNECTIVITY is due to Savitch [Sav]. Using the fact that S-T CONNECTIVITY (for *directed* graphs) is complete for nondeterministic logspace (\mathbf{NL}), this result is equivalent to the fact that $\mathbf{NL} \subset \mathbf{L}^c$, where \mathbf{L}^c is the class of languages that can be decided deterministic space $O(\log^c n)$. The latter formulation (and its generalization $\mathbf{NSPACE}(s(n)) \subset \mathbf{DSPACE}(s(n)^2)$) is known as Savitch’s Theorem. The randomized algorithm for UNDIRECTED S-T CONNECTIVITY was given by Aleliunas, Karp, Lipton, Lovász, and Rackoff [AKL⁺], and was recently derandomized by Reingold [Rei2] (see Section 4.4).

For more background on random walks, mixing time, and the Markov Chain Monte Carlo Method, we refer the reader to [MU, MR, Ran]. The use of this method for counting perfect matchings is due to [Bro, JS, JSV].

The bound on hitting time given in Theorem 2.49 is not tight; for example, it can be improved to $\Theta(n^2)$ for regular graphs that are simple (have no self-loops or parallel edges) [KLNS].

Even though we will focus primarily on undirected graphs (for example, in our study of expanders in Chapter 4), much of what we do generalizes to regular, or even Eulerian, digraphs. See Problem 2.10 and the references [Mih, Fil, RTV2]. Problem 2.10, Part 2 is from [Fil].

The Spectral Theorem (Thm. 2.54) can be found in any standard textbook on linear algebra. Problem 2.9 is from [Lov3]; Alon and Sudakov [AS2] strengthen it to show $\gamma(G) = \Omega(1/dDn)$, which implies a tight bound of $\gamma(G) = \Omega(1/n^2)$ for simple graphs (where $D = O(n/d)$).

Spectral Graph Theory is a rich subject, with many applications beyond the scope of this text; see the survey by Spielman [Spi] and references therein.

One significant omission from this chapter is the usefulness of randomness for *verifying proofs*. Recall that **NP** is the class of languages having membership proofs that can be verified in **P**. Thus it is natural to consider proof verification that is probabilistic, leading to the class **MA**, as well as a larger class **AM**, where the proof itself can depend on the randomness chosen by the verifier [BM]. (These are both subclasses of the class **IP** of languages having *interactive proof systems* [GMR1].) There are languages, such as GRAPH NONISOMORPHISM, that are in **AM** but are not known to be in **NP** [GMW]. “Derandomizing” these proof systems (e.g. proving **AM** = **NP**) would show that GRAPH NONISOMORPHISM is in **NP**, i.e. that there are short proofs that two graphs are nonisomorphic. Similarly to sampling and sublinear-time algorithms, randomized proof verification can also enable one to read only a small portion of an appropriately encoded **NP** proof, leading to the celebrated PCP Theorem and its applications to hardness of approximation [FGL⁺, AS1, ALM⁺]. For more about interactive proofs and PCPs, see [Vad, Gol6, AB].

3

Basic Derandomization Techniques

In the previous chapter, we saw some striking examples of the power of randomness for the design of efficient algorithms:

- POLYNOMIAL IDENTITY TESTING in **co-RP**.
- $[\times(1 + \varepsilon)]$ -APPROX #DNF in **prBPP**.
- PERFECT MATCHING in **RNC**.
- UNDIRECTED S-T CONNECTIVITY in **RL**.
- Approximating MAXCUT in probabilistic polynomial time.

This was of course only a small sample; there are entire texts on randomized algorithms. (See the notes and references for Chapter 2.)

In the rest of this survey, we will turn towards *derandomization* — trying to remove the randomness from these algorithms. We will achieve this for some of the specific algorithms we studied, and also consider the larger question of whether *all* efficient randomized algorithms can be derandomized, e.g. does **BPP = P**? **RL = L**? **RNC = NC**?

In this chapter, we will introduce a variety of “basic” derandomization techniques. These will each be deficient in that they are either infeasible (e.g. cannot be carried out in polynomial time) or specialized (e.g. apply only in very specific circumstances). But it will be useful to have these as tools before we proceed to study more sophisticated tools for derandomization (namely, the “pseudorandom objects” of Chapters 4–7).

3.1 Enumeration

We are interested in quantifying how much savings randomization provides. One way of doing this is to find the smallest possible upper bound on the deterministic time complexity of languages in **BPP**. For example, we would like to know which of the following complexity classes contain **BPP**:

Definition 3.1 (Deterministic Time Classes). ¹

$$\begin{aligned}
 \mathbf{DTIME}(t(n)) &= \{L : L \text{ can be decided deterministically in time } O(t(n))\} \\
 \mathbf{P} &= \cup_c \mathbf{DTIME}(n^c) && \text{("polynomial time")} \\
 \tilde{\mathbf{P}} &= \cup_c \mathbf{DTIME}(2^{(\log n)^c}) && \text{("quasipolynomial time")} \\
 \mathbf{SUBEXP} &= \cap_\epsilon \mathbf{DTIME}(2^{n^\epsilon}) && \text{("subexponential time")} \\
 \mathbf{EXP} &= \cup_c \mathbf{DTIME}(2^{n^c}) && \text{("exponential time")}
 \end{aligned}$$

The “Time Hierarchy Theorem” of complexity theory implies that all of these classes are distinct, i.e. $\mathbf{P} \subsetneq \tilde{\mathbf{P}} \subsetneq \mathbf{SUBEXP} \subsetneq \mathbf{EXP}$. More generally, it says that $\mathbf{DTIME}(o(t(n)/\log t(n))) \subsetneq \mathbf{DTIME}(t(n))$ for any efficiently computable time bound t . (What is difficult in complexity theory is separating classes that involve different computational resources, like deterministic time vs. nondeterministic time.)

Enumeration is a derandomization technique that enables us to deterministically simulate any randomized algorithm with an exponential slowdown.

Proposition 3.2. $\mathbf{BPP} \subset \mathbf{EXP}$.

Proof. If L is in \mathbf{BPP} , then there is a probabilistic polynomial-time algorithm A for L running in time $t(n)$ for some polynomial t . As in the proof of Theorem 2.30, we write $A(x; r)$ for A ’s output on input $x \in \{0, 1\}^n$ and coin tosses $r \in \{0, 1\}^{m(n)}$, where we may assume $m(n) \leq t(n)$ without loss of generality. Then:

$$\Pr[A(x; r) \text{ accepts}] = \frac{1}{2^{m(n)}} \sum_{r \in \{0, 1\}^{m(n)}} A(x; r)$$

We can compute the right-hand side of the above expression in deterministic time $2^{m(n)} \cdot t(n)$. \square

We see that the enumeration method is *general* in that it applies to all \mathbf{BPP} algorithms, but it is *infeasible* (taking exponential time). However, if the algorithm uses only a small number of random bits, it becomes feasible:

Proposition 3.3. If L has a probabilistic polynomial-time algorithm that runs in time $t(n)$ and uses $m(n)$ random bits, then $L \in \mathbf{DTIME}(t(n) \cdot 2^{m(n)})$. In particular, if $t(n) = \text{poly}(n)$ and $m(n) = O(\log n)$, then $L \in \mathbf{P}$.

Thus an approach to proving $\mathbf{BPP} = \mathbf{P}$ is to show that the number of random bits used by any \mathbf{BPP} algorithm can be reduced to $O(\log n)$. We will explore this approach in Chapter 7. However, to date, Proposition 3.2 remains the best unconditional upper-bound we have on the deterministic time-complexity of \mathbf{BPP} .

¹ Often $\mathbf{DTIME}(\cdot)$ is written as $\mathbf{TIME}(\cdot)$, but we include the \mathbf{D} to emphasize that it refers to deterministic rather than randomized algorithms. Also, in some contexts (e.g. cryptography), “exponential time” refers to $\mathbf{DTIME}(2^{O(n)})$ and “subexponential time” to be $\mathbf{DTIME}(2^{o(n)})$; we will use \mathbf{E} to denote the former class when it arises in Chapter 7.

Open Problem 3.4. Is **BPP** “closer” to **P** or **EXP**? Is $\mathbf{BPP} \subset \tilde{\mathbf{P}}$? Is $\mathbf{BPP} \subset \mathbf{SUBEXP}$?

3.2 Nonconstructive/Nonuniform Derandomization

Next we look at a derandomization technique that can be implemented efficiently but requires some nonconstructive “advice” that depends on the input length.

Proposition 3.5. If $A(x; r)$ is a randomized algorithm for a language L that has error probability smaller than 2^{-n} on inputs x of length n , then for every n , there exists a fixed sequence of coin tosses r_n such that $A(x; r_n)$ is correct for all $x \in \{0, 1\}^n$.

Proof. We use the Probabilistic Method. Consider R_n chosen uniformly at random from $\{0, 1\}^{r(n)}$, where $r(n)$ is the number of coin tosses used by A on inputs of length n . Then

$$\begin{aligned} \Pr[\exists x \in \{0, 1\}^n \text{ s.t. } A(x; R_n) \text{ incorrect on } x] &\leq \sum_x \Pr[A(x; R_n) \text{ incorrect on } x] \\ &< 2^n \cdot 2^{-n} = 1 \end{aligned}$$

Thus, there exists a fixed value $R_n = r_n$ that yields a correct answer for all $x \in \{0, 1\}^n$. \square

The advantage of this method over enumeration is that once we have the fixed string r_n , computing $A(x; r_n)$ can be done in polynomial time. However, the proof that r_n exists is nonconstructive; it is not clear how to find it in less than exponential time.

Note that we know that we can reduce the error probability of any **BPP** (or **RP**, **RL**, **RNC**, etc.) algorithm to smaller than 2^{-n} by repetitions, so this proposition is always applicable. However, we begin by looking at some interesting special cases.

Example 3.6 (Perfect Matching). We apply the proposition to Algorithm 2.7. Let $G = (V, E)$ be a bipartite graph with m vertices on each side, and let $A^G(x_{1,1}, \dots, x_{m,m})$ be the matrix that has entries $A_{i,j}^G = x_{i,j}$ if $(i, j) \in E$, and $A_{i,j}^G = 0$ if $(i, j) \notin E$. Recall that the polynomial $\det(A^G(x))$ is nonzero if and only if G has a perfect matching. Let $S_m = \{0, 1, 2, \dots, m2^{m^2}\}$. We argued that, by the Schwartz–Zippel Lemma, if we choose $\alpha \stackrel{R}{\leftarrow} S_m^{m^2}$ at random and evaluate $\det(A^G(\alpha))$, we can determine whether $\det(A^G(x))$ is zero with error probability at most $m/|S|$ which is smaller than 2^{-m^2} . Since a bipartite graph with m vertices per side is specified by a string of length $n = m^2$, by Proposition 3.5 we know that for every m , there exists an $\alpha_m \in S_m^{m^2}$ such that $\det(A^G(\alpha_m)) \neq 0$ if and only if G has a perfect matching, for every bipartite graph G with m vertices on each side.

Open Problem 3.7. Can we find such an $\alpha_m \in \{0, \dots, m2^{m^2}\}^{m^2}$ explicitly, i.e., *deterministically* and *efficiently*? An **NC** algorithm (i.e. parallel time $\text{polylog}(m)$ with $\text{poly}(m)$ processors) would put PERFECT MATCHING in deterministic **NC**, but even a subexponential-time algorithm would be interesting.

Example 3.8 (Universal Traversal Sequences). Let G be a connected d -regular undirected multigraph on n vertices. From Theorem 2.49, we know that a random walk of $\text{poly}(n, d)$ steps from any start vertex will visit any other vertex with high probability. By increasing the length of the walk by a polynomial factor, we can ensure that *every* vertex is visited with probability greater than $1 - 2^{-nd \log n}$. By the same reasoning as in the previous example, we conclude that for every pair (n, d) , there exists a *universal traversal sequence* $w \in \{1, 2, \dots, d\}^{\text{poly}(n, d)}$ such that for every n -vertex, d -regular, connected G and every vertex s in G , if we start from s and follow w then we will visit the entire graph.

Open Problem 3.9. Can we construct such a universal traversal sequence explicitly (e.g. in polynomial time or even logarithmic space)?

There has been substantial progress towards resolving this question in the positive; see Section 4.4.

We now cast the nonconstructive derandomizations provided by Proposition 3.5 in the language of “nonuniform” complexity classes.

Definition 3.10. Let \mathbf{C} be a class of languages, and $\mathbf{a} : \mathbb{N} \rightarrow \mathbb{N}$ be a function. Then \mathbf{C}/\mathbf{a} is the class of languages defined as follows: $L \in \mathbf{C}/\mathbf{a}$ if there exists $L' \in \mathbf{C}$, and $\alpha_1, \alpha_2, \dots \in \{0, 1\}^*$ with $|\alpha_n| \leq \mathbf{a}(n)$, such that $x \in L \Leftrightarrow (x, \alpha_{|x|}) \in L'$. The α 's are called the *advice strings*.

\mathbf{P}/\mathbf{poly} is the class $\bigcup_c \mathbf{P}/n^c$, i.e. polynomial time with polynomial advice.

A basic result in complexity theory is that \mathbf{P}/\mathbf{poly} is exactly the class of languages that can be decided by polynomial-sized Boolean circuits:

Fact 3.11. $L \in \mathbf{P}/\mathbf{poly}$ iff there is a sequence of Boolean circuits $\{C_n\}_{n \in \mathbb{N}}$ and a polynomial p such that for all n

- (1) $C_n : \{0, 1\}^n \rightarrow \{0, 1\}$ decides $L \cap \{0, 1\}^n$
 - (2) $|C_n| \leq p(n)$.
-

We refer to \mathbf{P}/\mathbf{poly} as a “nonuniform” model of computation because it allows for different, unrelated “programs” for each input length (e.g. the circuits C_n , or the advice α_n), in contrast to classes like \mathbf{P} , \mathbf{BPP} , and \mathbf{NP} , that require a single “program” of constant size specifying how the computation should behave for inputs of arbitrary length. Although \mathbf{P}/\mathbf{poly} contains some undecidable problems,² people generally believe that $\mathbf{NP} \not\subseteq \mathbf{P}/\mathbf{poly}$. Indeed, trying to prove lower bounds on circuit size is one of the main approaches to proving $\mathbf{P} \neq \mathbf{NP}$, since circuits seem much more concrete and combinatorial than Turing machines. However, this too has turned out to be quite difficult; the best circuit lower bound known for computing an explicit function is roughly $5n$.

Proposition 3.5 directly implies:

²Consider the unary version of the HALTING PROBLEM, which can be decided in constant time given advice $\alpha_n \in \{0, 1\}$ that specifies whether the n 'th Turing machine halts or not.

Corollary 3.12. $\mathbf{BPP} \subset \mathbf{P}/\text{poly}$.

A more general meta-theorem is that “nonuniformity is more powerful than randomness.”

3.3 Nondeterminism

Although physically unrealistic, nondeterminism has proven to be a very useful resource in the study of computational complexity (e.g. leading to the class \mathbf{NP}). Thus it is natural to study how it compares in power to randomness. Intuitively, with nondeterminism we should be able to guess a “good” sequence of coin tosses for a randomized algorithm and then do the computation deterministically. This intuition does apply directly for randomized algorithms with 1-sided error:

Proposition 3.13. $\mathbf{RP} \subset \mathbf{NP}$.

Proof. Let $L \in \mathbf{RP}$ and A be a randomized algorithm that decides it. A polynomial-time verifiable witness that $x \in L$ is any sequence of coin tosses r such that $A(x; r) = \text{accept}$. \square

However, for 2-sided error (\mathbf{BPP}), containment in \mathbf{NP} is not clear. Even if we guess a ‘good’ random string (one that leads to a correct answer), it is not clear how we can verify it in polynomial time. Indeed, it is consistent with current knowledge that \mathbf{BPP} equals \mathbf{NEXP} (nondeterministic exponential time)! Nevertheless, there is a sense in which we can show that \mathbf{BPP} is no more powerful than \mathbf{NP} :

Theorem 3.14. If $\mathbf{P} = \mathbf{NP}$, then $\mathbf{P} = \mathbf{BPP}$.

Proof. For any language $L \in \mathbf{BPP}$, we will show how to express membership in L using two quantifiers. That is, for some polynomial-time predicate P ,

$$x \in L \iff \exists y \forall z P(x, y, z), \tag{3.1}$$

where we quantify over y and z of length $\text{poly}(|x|)$.

Assuming $\mathbf{P} = \mathbf{NP}$, we can replace $\forall z P(x, y, z)$ by a polynomial-time predicate $Q(x, y)$, because the language $\{(x, y) : \forall z P(x, y, z)\}$ is in $\mathbf{co-NP} = \mathbf{P}$. Then $L = \{x : \exists y Q(x, y)\} \in \mathbf{NP} = \mathbf{P}$.

To obtain the two-quantifier expression (3.1), consider a randomized algorithm A for L , and assume w.l.o.g. that its error probability is smaller than 2^{-n} and that it uses $m = \text{poly}(n)$ coin tosses. Let $A_x \subset \{0, 1\}^m$ be the set of coin tosses r for which $A(x; r) = 1$. We will show that if x is in L , there exist m “shifts” (or “translations”) of A_x that cover all points in $\{0, 1\}^m$. (Notice that this is a $\exists \forall$ statement.) Intuitively, this should be possible because A_x contains all but an exponentially small fraction of $\{0, 1\}^m$. On the other hand if $x \notin L$, then no m shifts of A_x can cover all of $\{0, 1\}^m$. Intuitively, this is because A_x is an exponentially small fraction of $\{0, 1\}^m$.

Formally, by a “shift of A_x ” we mean a set of the form $A_x \oplus s = \{r \oplus s : r \in A_x\}$ for some

$s \in \{0, 1\}^m$; note that $|A_x \oplus s| = |A_x|$. We will show

$$\begin{aligned}
x \in L &\Rightarrow \exists s_1, s_2, \dots, s_m \in \{0, 1\}^m \forall r \in \{0, 1\}^m \quad r \in \bigcup_{i=1}^m (A_x \oplus s_i) \\
&\Leftrightarrow \exists s_1, s_2, \dots, s_m \in \{0, 1\}^m \forall r \in \{0, 1\}^m \quad \bigvee_{i=1}^m (A(x; r \oplus s_i) = 1); \\
x \notin L &\Rightarrow \forall s_1, s_2, \dots, s_m \in \{0, 1\}^m \exists r \in \{0, 1\}^m \quad r \notin \bigcup_{i=1}^m (A_x \oplus s_i) \\
&\Leftrightarrow \forall s_1, s_2, \dots, s_m \in \{0, 1\}^m \exists r \in \{0, 1\}^m \quad \neg \bigvee_{i=1}^m (A(x; r \oplus s_i) = 1).
\end{aligned}$$

We prove both parts by the Probabilistic Method, starting with the second (which is simpler).

$x \notin L$ Let s_1, \dots, s_m be arbitrary, and choose $R \stackrel{\text{R}}{\leftarrow} \{0, 1\}^m$ at random. Now A_x and hence each $A_x \oplus s_i$ contains less than a 2^{-n} fraction of $\{0, 1\}^m$. So, by a union bound,

$$\begin{aligned}
\Pr[R \in \bigcup_i (A_x \oplus s_i)] &\leq \sum_i \Pr[R \in A_x \oplus s_i] \\
&< m \cdot 2^{-n} < 1,
\end{aligned}$$

for sufficiently large n . In particular, there exists an $r \in \{0, 1\}^m$ such that $r \notin \bigcup_i (A_x \oplus s_i)$.

$x \in L$: Choose $S_1, S_2, \dots, S_m \stackrel{\text{R}}{\leftarrow} \{0, 1\}^m$. Then, for every fixed r , we have:

$$\begin{aligned}
\Pr[r \notin \bigcup_i (A_x \oplus S_i)] &= \prod_i \Pr[r \notin A_x \oplus S_i] \\
&= \prod_i \Pr[S_i \notin A_x \oplus r] \\
&< (2^{-n})^m,
\end{aligned}$$

since A_x and hence $A_x \oplus r$ contains more than a $1 - 2^{-n}$ fraction of $\{0, 1\}^m$. By a union bound, we have:

$$\Pr[\exists r \quad r \notin \bigcup_i (A_x \oplus S_i)] < 2^m \cdot (2^{-n})^m \leq 1.$$

Thus, there exist s_1, \dots, s_m such that $\bigcup_i (A_x \oplus s_i)$ contains all points r in $\{0, 1\}^m$. □

Readers familiar with complexity theory will recognize the above proof as showing that **BPP** is contained in the 2nd level of the *polynomial-time hierarchy* (**PH**). In general, the k 'th level of the **PH** contains all languages that satisfy a k -quantifier expression analogous to (3.1).

3.4 The Method of Conditional Expectations

In the previous sections, we saw several derandomization techniques (enumeration, nonuniformity, nondeterminism) that are general in the sense that they apply to all of **BPP**, but are infeasible

Note that to implement this method, we need to compute $P(r_1, r_2, \dots, r_i)$ deterministically, and this may be infeasible. However, there are nontrivial algorithms where this method does work, often for *search* problems rather than decision problems, and where we measure not a boolean outcome (e.g. whether A is correct as above) but some other measure of quality of the output. Below we see one such example, where it turns out to yield a natural “greedy algorithm”.

3.4.2 Derandomized MAXCUT Approximation

Recall the MAXCUT problem:

Computational Problem 3.15 (Computational Problem 2.38, rephrased).

MAXCUT: Given a graph $G = (V, E)$, find a partition S, T of V (i.e. $S \cup T = V$, $S \cap T = \emptyset$) maximizing the size of the set $\text{cut}(S, T) = \{\{u, v\} \in E : u \in S, v \in T\}$.

We saw a simple randomized algorithm that finds a cut of (expected) size at least $|E|/2$ (not counting any self-loops, which can never be cut), which we now phrase in a way suitable for derandomization.

Algorithm 3.16 (randomized MAXCUT, rephrased).

Input: a graph $G = ([N], E)$ (with no self-loops)

Flip N coins r_1, r_2, \dots, r_N , put vertex i in S if $r_i = 0$ and in T if $r_i = 1$. Output (S, T) .

To derandomize this algorithm using the Method of Conditional Expectations, define the conditional expectation

$$e(r_1, r_2, \dots, r_i) \stackrel{\text{def}}{=} \mathbb{E}_{R_1, R_2, \dots, R_N} \left[|\text{cut}(S, T)| \mid R_1 = r_1, R_2 = r_2, \dots, R_i = r_i \right]$$

to be the expected cut size when the random choices for the first i coins are fixed to r_1, r_2, \dots, r_i .

We know that when no random bits are fixed, $e[\Lambda] = |E|/2$ (because each edge is cut with probability $1/2$), and all we need to calculate is $e(r_1, r_2, \dots, r_i)$ for $1 \leq i \leq N$. For this particular algorithm it turns out that the quantity is not hard to compute. Let $S_i \stackrel{\text{def}}{=} \{j : j \leq i, r_j = 0\}$ (resp. $T_i \stackrel{\text{def}}{=} \{j : j \leq i, r_j = 1\}$) be the set of vertices in S (resp. T) after we determine r_1, \dots, r_i , and $U_i \stackrel{\text{def}}{=} \{i+1, i+2, \dots, N\}$ be the “undecided” vertices that have not been put into S or T . Then

$$e(r_1, r_2, \dots, r_i) = |\text{cut}(S_i, T_i)| + 1/2 (|\text{cut}(U_i, [N])|). \tag{3.2}$$

Note that $\text{cut}(U_i, [N])$ is the set of *unordered* edges that have at least one endpoint in U_i . Now we can deterministically select a value for r_{i+1} , by computing and comparing $e(r_1, r_2, \dots, r_i, 0)$ and $e(r_1, r_2, \dots, r_i, 1)$.

In fact, the decision on r_{i+1} can be made even simpler than computing (3.2) in its entirety, by observing that the set $\text{cut}(U_{i+1}, [N])$ is independent of the choice of r_{i+1} . Therefore, to maximize $e(r_1, r_2, \dots, r_i, r_{i+1})$, it is enough to choose r_{i+1} that maximizes the $|\text{cut}(S, T)|$ term. This term increases by either $|\text{cut}(\{i+1\}, T_i)|$ or $|\text{cut}(\{i+1\}, S_i)|$ depending on whether we place vertex $i+1$ in S or T , respectively. To summarize, we have

$$e(r_1, \dots, r_i, 0) - e(r_1, \dots, r_i, 1) = |\text{cut}(\{i+1\}, T_i)| - |\text{cut}(\{i+1\}, S_i)|.$$

This gives rise to the following deterministic algorithm, which is guaranteed to always find a cut of size at least $|E|/2$:

Algorithm 3.17 (deterministic MAXCUT approximation).

Input: a graph $G = ([N], E)$ (with no self-loops)

- (1) Set $S = \emptyset, T = \emptyset$
 - (2) For $i = 0, \dots, N - 1$:
 - (a) If $|\text{cut}(\{i + 1\}, T)| > |\text{cut}(\{i + 1\}, S)|$, set $S \leftarrow S \cup \{i + 1\}$,
 - (b) Else set $T \leftarrow T \cup \{i + 1\}$.
-

Note that this is the natural “greedy” algorithm for this problem. In other cases, the Method of Conditional Expectations yields algorithms that, while still arguably “greedy,” would have been much less easy to find directly. Thus, designing a randomized algorithm and then trying to derandomize it can be a useful paradigm for the design of deterministic algorithms even if the randomization does not provide gains in efficiency.

3.5 Pairwise Independence

3.5.1 An Example

As a motivating example for pairwise independence, we give another way of derandomizing the MAXCUT approximation algorithm discussed above. Recall the analysis of the randomized algorithm:

$$E[|\text{cut}(S)|] = \sum_{(i,j) \in E} \Pr[R_i \neq R_j] = |E|/2,$$

where R_1, \dots, R_N are the random bits of the algorithm. The key observation is that this analysis applies for any distribution on (R_1, \dots, R_N) satisfying $\Pr[R_i \neq R_j] = 1/2$ for each $i \neq j$. Thus, they do not need to be completely independent random variables; it suffices for them to be *pairwise independent*. That is, each R_i is an unbiased random bit, and for each $i \neq j$, R_i is independent from R_j .

This leads to the question: Can we generate N pairwise independent bits using less than N truly random bits? The answer turns out to be *yes*, as illustrated by the following construction.

Construction 3.18 (pairwise independent bits). Let B_1, \dots, B_k be k independent unbiased random bits. For each nonempty $S \subset [k]$, let R_S be the random variable $\bigoplus_{i \in S} B_i$, where \oplus denotes XOR.

Proposition 3.19. The $2^k - 1$ random variables R_S in Construction 3.18 are pairwise independent unbiased random bits.

Proof. It is evident that each R_S is unbiased. For pairwise independence, consider any two nonempty sets $S \neq T \subset [k]$. Then:

$$\begin{aligned} R_S &= R_{S \cap T} \oplus R_{S \setminus T} \\ R_T &= R_{S \cap T} \oplus R_{T \setminus S}. \end{aligned}$$

Note that $R_{S \cap T}$, $R_{S \setminus T}$ and $R_{T \setminus S}$ are independent as they depend on disjoint subsets of the B_i 's, and at least two of these subsets are nonempty. This implies that (R_S, R_T) takes each value in $\{0, 1\}^2$ with probability $1/4$. \square

Note that this gives us a way to generate N pairwise independent bits from $\lceil \log(N + 1) \rceil$ independent random bits. Thus, we can reduce the randomness required by the MAXCUT algorithm to logarithmic, and then we can obtain a deterministic algorithm by enumeration.

Algorithm 3.20 (deterministic MAXCUT algorithm II).

Input: a graph $G = ([N], E)$ (with no self-loops)

For *all* sequences of bits $b_1, b_2, \dots, b_{\lceil \log(N+1) \rceil}$, run the randomized MAXCUT algorithm using coin tosses $(r_S = \bigoplus_{i \in S} b_i)_{S \neq \emptyset}$ and choose the largest cut thus obtained.

Since there are at most $2(N + 1)$ sequences of b_i 's, the derandomized algorithm still runs in polynomial time. It is slower than the algorithm we obtained by the Method of Conditional Expectations (Algorithm 3.17), but it has the advantage of using logarithmic workspace and being parallelizable. The derandomization can be sped up using almost pairwise independence (at the price of a slightly worse approximation factor); see Problem 3.4.

3.5.2 Pairwise Independent Hash Functions

Some applications require pairwise independent random variables that take values from a larger range, e.g. we want $N = 2^n$ pairwise independent random variables, each of which is uniformly distributed in $\{0, 1\}^m = [M]$. (Throughout this survey, we will often use the convention that a lowercase letter equals the logarithm (base 2) of the corresponding capital letter.) The naive approach is to repeat the above algorithm for the individual bits m times. This uses roughly $n \cdot m = (\log M)(\log N)$ initial random bits, which is no longer logarithmic in N if M is nonconstant. Below we will see that we can do much better. But first some definitions.

A sequence of N random variables each taking a value in $[M]$ can be viewed as a distribution on sequences in $[M]^N$. Another interpretation of such a sequence is as a mapping $f : [N] \rightarrow [M]$. The latter interpretation turns out to be more useful when discussing the computational complexity of the constructions.

Definition 3.21 (pairwise independent hash functions). A family (i.e. multiset) of functions $\mathcal{H} = \{h : [N] \rightarrow [M]\}$ is *pairwise independent* if the following two conditions hold when $H \stackrel{R}{\leftarrow} \mathcal{H}$ is a function chosen uniformly at random from \mathcal{H} :

- (1) $\forall x \in [N]$, the random variable $H(x)$ is uniformly distributed in $[M]$.

(2) $\forall x_1 \neq x_2 \in [N]$, the random variables $H(x_1)$ and $H(x_2)$ are independent.

Equivalently, we can combine the two conditions and require that

$$\forall x_1 \neq x_2 \in [N], \forall y_1, y_2 \in [M], \Pr_{H \leftarrow \mathcal{H}} [H(x_1) = y_1 \wedge H(x_2) = y_2] = \frac{1}{M^2}.$$

Note that the probability above is over the random choice of a function from the family \mathcal{H} . This is why we talk about a family of functions rather than a single function. The description in terms of functions makes it natural to impose a strong efficiency requirement:

Definition 3.22. A family of functions $\mathcal{H} = \{h : [N] \rightarrow [M]\}$ is *explicit* if given the description of h and $x \in [N]$, the value $h(x)$ can be computed in time $\text{poly}(\log N, \log M)$.

Pairwise independent hash functions are sometimes referred to as *strongly 2-universal hash functions*, to contrast with the weaker notion of *2-universal hash functions*, which requires only that $\Pr[H(x_1) = H(x_2)] \leq 1/M$ for all $x_1 \neq x_2$. (Note that this property is all we needed for the deterministic MAXCUT algorithm, and it allows for a small savings in that we can also include $S = \emptyset$ in Construction 3.18.)

Below we present another construction of a pairwise independent family.

Construction 3.23 (pairwise independent hash functions from linear maps). Let \mathbb{F} be a finite field. Define the family of functions $\mathcal{H} = \{h_{a,b} : \mathbb{F} \rightarrow \mathbb{F}\}_{a,b \in \mathbb{F}}$ where $h_{a,b}(x) = ax + b$.

Proposition 3.24. The family of functions \mathcal{H} in Construction 3.23 is pairwise independent.

Proof. Notice that the graph of the function $h_{a,b}(x)$ is the line with slope a and y -intercept b . Given $x_1 \neq x_2$ and y_1, y_2 , there is exactly one such line containing the points (x_1, y_1) and (x_2, y_2) (namely, the line with slope $a = (y_1 - y_2)/(x_1 - x_2)$ and y -intercept $b = y_1 - ax_1$). Thus, the probability over a, b that $h_{a,b}(x_1) = y_1$ and $h_{a,b}(x_2) = y_2$ equals the reciprocal of the number of lines, namely $1/|\mathbb{F}|^2$. \square

This construction uses $2 \log |\mathbb{F}|$ random bits, since we have to choose a and b at random from \mathbb{F} to get a function $h_{a,b} \leftarrow \mathcal{H}$. Compare this to $|\mathbb{F}| \log |\mathbb{F}|$ bits required to choose a truly random function, and $(\log |\mathbb{F}|)^2$ bits for repeating the construction of Proposition 3.19 for each output bit.

Note that evaluating the functions of Construction 3.23 requires a description of the field \mathbb{F} that enables us to perform addition and multiplication of field elements. Thus we take a brief aside to review the complexity of constructing and computing in finite fields.

Remark 3.25 (constructing finite fields). Recall that for every prime power $q = p^k$ there is a field \mathbb{F}_q (often denoted $\text{GF}(q)$) of size q , and this field is unique up to isomorphism (renaming elements). The prime p is called the *characteristic* of the field. \mathbb{F}_q has a description of length $O(\log q)$ enabling addition, multiplication, and division to be performed in polynomial time (i.e.

time $\text{poly}(\log q)$. (This description is simply an irreducible polynomial f of degree k over $\mathbb{F}_p = \mathbb{Z}_p$.) To construct this description, we first need to determine the characteristic p ; finding a prime p of a desired bitlength ℓ can be done probabilistically in time $\text{poly}(\ell) = \text{poly}(\log p)$ and deterministically in time $\text{poly}(2^\ell) = \text{poly}(p)$. Then given p and k , a description of \mathbb{F}_q (for $q = p^k$) can be found probabilistically in time $\text{poly}(\log p, k) = \text{poly}(\log q)$ and deterministically in time $\text{poly}(p, k)$. Note that for both steps, the deterministic algorithms are exponential in the bitlength of the characteristic p . Thus, for computational purposes, a convenient choice is often to instead take $p = 2$ and k large, in which case everything can be done deterministically in time $\text{poly}(k) = \text{poly}(\log q)$.

Using a finite fields of size 2^n as suggested above, we obtain an explicit construction of pairwise independent hash functions $\mathcal{H}_{n,n} = \{h : \{0, 1\}^n \rightarrow \{0, 1\}^n\}$ for every n . What if we want a family $\mathcal{H}_{n,m}$ of pairwise independent hash functions where the input length n and output length m are not equal? For $n < m$, we can take hash functions h from $\mathcal{H}_{m,m}$ and restrict their domain to $\{0, 1\}^n$ by defining $h'(x) = h(x \circ 0^{m-n})$. In the case that $m < n$, we can take h from $\mathcal{H}_{n,n}$ and throw away $n - m$ bits of the output. That is, define $h'(x) = h(x)|_m$, where $h(x)|_m$ denotes the first m bits of $h(x)$.

In both cases, we use $2 \max\{n, m\}$ random bits. This is the best possible when $m \geq n$. When $m < n$, it can be reduced to $m + n$ random bits by using $(ax)|_m + b$ where $b \in \{0, 1\}^m$ instead of $(ax + b)|_m$. Summarizing:

Theorem 3.26. For every $n, m \in \mathbb{N}$, there is an explicit family of pairwise independent functions $\mathcal{H}_{n,m} = \{h : \{0, 1\}^n \rightarrow \{0, 1\}^m\}$ where a random function from $\mathcal{H}_{n,m}$ can be selected using $\max\{m, n\} + m$ random bits.

Problem 3.5 shows that $\max\{m, n\} + m$ random bits is essentially optimal.

3.5.3 Hash Tables

The original motivating application for pairwise independent functions was for hash tables. Suppose we want to store a set $S \subset [N]$ of values and answer queries of the form “Is $x \in S$?” efficiently (or, more generally, acquire some piece of data associated with key x in case $x \in S$). A simple solution is to have a table T such that $T[x] = 1$ if and only if $x \in S$. But this requires N bits of storage, which is inefficient if $|S| \ll N$.

A better solution is to use hashing. Assume that we have a hash function $h : [N] \rightarrow [M]$ for some M to be determined later. Let T be a table of size M , initially empty. For each $x \in [N]$, we let $T[h(x)] = x$ if $x \in S$. So to test whether a given $y \in S$, we compute $h(y)$ and check if $T[h(y)] = y$. In order for this construction to be well-defined, we need h to be one-to-one on the set S . Suppose we choose a random function H from $[N]$ to $[M]$. Then, for any set S of size at most K , the probability that there are any collisions is

$$\Pr[\exists x \neq y \text{ s.t. } H(x) = H(y)] \leq \sum_{x \neq y \in S} \Pr[H(x) = H(y)] \leq \binom{K}{2} \cdot \frac{1}{M} < \varepsilon$$

for $M = K^2/\varepsilon$. Notice that the above analysis does not require H to be a completely random function; it suffices that H be pairwise independent (or even 2-universal). Thus using Theorem 3.26,

we can generate and store H using $O(\log N)$ random bits. The storage required for the table T is $O(M \log N) = O(K^2 \log N)$ bits, which is much smaller than N when $K = N^{o(1)}$. Note that to uniquely represent a set of size K , we need space at least $\log \binom{N}{K} = \Omega(K \log N)$ (when $K \leq N^{.99}$). In fact, there is a data structure achieving a matching space upper bound of $O(K \log N)$, which works by taking $M = O(K)$ in the above construction and using additional hash functions to separate the (few) collisions that will occur.

Often, when people analyze applications of hashing in computer science, they model the hash function as a truly random function. However, the domain of the hash function is often exponentially large, and thus it is infeasible to even write down a truly random hash function. Thus, it would be preferable to show that some explicit family of hash function works for the application with similar performance. In many cases (such as the one above), it can be shown that pairwise independence (or k -wise independence, as discussed below) suffices.

3.5.4 Randomness-Efficient Error Reduction and Sampling

Suppose we have a **BPP** algorithm for a language L that has a constant error probability. We want to reduce the error to 2^{-k} . We have already seen that this can be done using $O(k)$ independent repetitions (by a Chernoff Bound). If the algorithm originally used m random bits, then we use $O(km)$ random bits after error reduction. Here we will see how to reduce the number of random bits required for error reduction by doing repetitions that are only pairwise independent.

To analyze this, we will need an analogue of the Chernoff Bound (Theorem 2.21) that applies to averages of pairwise independent random variables. This will follow from Chebyshev's Inequality, which bounds the deviations of a random variable X from its mean μ in terms its *variance* $\text{Var}[X] = \text{E}[(X - \mu)^2] = \text{E}[X^2] - \mu^2$.

Lemma 3.27 (Chebyshev's Inequality). If X is a random variable with expectation μ , then

$$\Pr[|X - \mu| \geq \varepsilon] \leq \frac{\text{Var}[X]}{\varepsilon^2}$$

Proof. Applying Markov's Inequality (Lemma 2.20) to the random variable $Y = (X - \mu)^2$, we have:

$$\Pr[|X - \mu| \geq \varepsilon] = \Pr[(X - \mu)^2 \geq \varepsilon^2] \leq \frac{\text{E}[(X - \mu)^2]}{\varepsilon^2} = \frac{\text{Var}[X]}{\varepsilon^2}.$$

□

We now use this to show that a sum of pairwise independent random variables is concentrated around its expectation.

Proposition 3.28 (Pairwise Independent Tail Inequality). Let X_1, \dots, X_t be pairwise independent random variables taking values in the interval $[0, 1]$, let $X = (\sum_i X_i)/t$, and $\mu = \text{E}[X]$. Then

$$\Pr[|X - \mu| \geq \varepsilon] \leq \frac{1}{t\varepsilon^2}.$$

Proof. Let $\mu_i = E[X_i]$. Then

$$\begin{aligned}
 \text{Var}[X] &= E[(X - \mu)^2] \\
 &= \frac{1}{t^2} \cdot E\left[\left(\sum_i (X_i - \mu_i)\right)^2\right] \\
 &= \frac{1}{t^2} \cdot \sum_{i,j} E[(X_i - \mu_i)(X_j - \mu_j)] \\
 &= \frac{1}{t^2} \cdot \sum_i E[(X_i - \mu_i)^2] \quad (\text{by pairwise independence}) \\
 &= \frac{1}{t^2} \cdot \sum_i \text{Var}[X_i] \\
 &\leq \frac{1}{t}
 \end{aligned}$$

Now apply Chebyshev's Inequality. □

While this requires less independence than the Chernoff Bound, notice that the error probability decreases only linearly rather than exponentially with the number t of samples.

Error Reduction. Proposition 3.28 tells us that if we use $t = O(2^k)$ pairwise independent repetitions, we can reduce the error probability of a **BPP** algorithm from $1/3$ to 2^{-k} . If the original **BPP** algorithm uses m random bits, then we can do this by choosing $h : \{0, 1\}^{k+O(1)} \rightarrow \{0, 1\}^m$ at random from a pairwise independent family, and running the algorithm using coin tosses $h(x)$ for all $x \in \{0, 1\}^{k+O(1)}$. This requires $m + \max\{m, k + O(1)\} = O(m + k)$ random bits.

	Number of Repetitions	Number of Random Bits
Independent Repetitions	$O(k)$	$O(km)$
Pairwise Independent Repetitions	$O(2^k)$	$O(m + k)$

Note that we saved substantially on the number of random bits, but paid a lot in the number of repetitions needed. To maintain a polynomial-time algorithm, we can only afford $k = O(\log n)$. This setting implies that if we have a **BPP** algorithm with constant error that uses m random bits, we have another **BPP** algorithm that uses $O(m + \log n) = O(m)$ random bits and has error $1/\text{poly}(n)$. That is, we can go from constant to inverse-polynomial error only paying a constant factor in randomness. (In fact, it turns out there is a way to achieve this with *no* penalty in randomness; see Problem 4.6.)

Sampling. Recall the SAMPLING problem: Given oracle access to a function $f : \{0, 1\}^m \rightarrow [0, 1]$, we want to approximate $\mu(f)$ to within an additive error of ε .

In Section 2.3.1, we saw that we can solve this problem with probability $1 - \delta$ by outputting the average of f on a random sample of $t = O(\log(1/\delta)/\varepsilon^2)$ points in $\{0, 1\}^m$, where the correctness follows from the Chernoff Bound. To reduce the number of truly random bits used, we can use a pairwise independent sample instead. Specifically, taking $t = 1/(\varepsilon^2\delta)$ pairwise independent points, we get an error probability of at most δ (by Proposition 3.28). To generate t pairwise independent

samples of m bits each, we need $O(m + \log t) = O(m + \log(1/\varepsilon) + \log(1/\delta))$ truly random bits.

	Number of Samples	Number of Random Bits
Truly Random Sample	$O((1/\varepsilon^2) \cdot \log(1/\delta))$	$O(m \cdot (1/\varepsilon^2) \cdot \log(1/\delta))$
Pairwise Independent Repetitions	$O((1/\varepsilon^2) \cdot (1/\delta))$	$O(m + \log(1/\varepsilon) + \log(1/\delta))$

Both of these sampling algorithms have a natural restricted structure. First, they choose all of their queries to the oracle f nonadaptively, based solely on their coin tosses and not based on the answers to previous queries. Second, their output is simply the average of the queried values, whereas the original sampling problem does not constrain the output function. It is useful to abstract these properties as follows.

Definition 3.29. A *sampler* Samp for domain size M is given “coin tosses” $x \stackrel{\text{R}}{\leftarrow} [N]$ and outputs a sequence of samples $z_1, \dots, z_t \in [M]$. We say that $\text{Samp} : [N] \rightarrow [M]^t$ is a (δ, ε) *averaging sampler* if for every function $f : [M] \rightarrow [0, 1]$, we have

$$\Pr_{(z_1, \dots, z_t) \stackrel{\text{R}}{\leftarrow} \text{Samp}(U_{[N]})} \left[\frac{1}{t} \sum_{i=1}^t f(z_i) > \mu(f) + \varepsilon \right] \leq \delta. \quad (3.3)$$

If Inequality 3.3 only holds for f with (boolean) range $\{0, 1\}$, we call Samp a *boolean averaging sampler*. We say that Samp is *explicit* if given $x \in [N]$ and $i \in [t]$, $\text{Samp}(x)_i$ can be computed in time $\text{poly}(\log N, \log t)$.

We note that, in contrast to the Chernoff Bound (Theorem 2.21) and the Pairwise Independent Tail Inequality (Proposition 3.28), this definition seems to only provide an error guarantee in one direction, namely that the sample average does not significantly *exceed* the global average (except with small probability). However, a guarantee in the other direction also follows by considering the function $1 - f$. Thus, up to a factor of 2 in the failure probability δ , the above definition is equivalent to requiring that $\Pr[|(1/t) \cdot \sum_i f(z_i) - \mu(f)| > \varepsilon] \leq \delta$. We choose to use a one-sided guarantee because it will make the connection to list-decodable codes (in Chapter 5) slightly cleaner.

Our pairwise-independent sampling algorithm can now be described as follows:

Theorem 3.30 (Pairwise Independent Sampler). For every $m \in \mathbb{N}$ and $\delta, \varepsilon \in [0, 1]$, there is an explicit (δ, ε) averaging sampler $\text{Samp} : \{0, 1\}^n \rightarrow (\{0, 1\}^m)^t$ using $n = O(m + \log(1/\varepsilon) + \log(1/\delta))$ random bits and $t = O(1/(\varepsilon^2 \delta))$ samples.

As we will see in subsequent chapters, averaging samplers are intimately related to the other pseudorandom objects we are studying (especially randomness extractors). In addition, some applications of samplers require samplers of this restricted form.

3.5.5 k -wise Independence

Our definition and construction of pairwise independent functions generalize naturally to k -wise independence for any k .

Definition 3.31 (*k-wise independent hash functions*). For $N, M, k \in \mathbb{N}$ such that $k \leq N$, a family of functions $\mathcal{H} = \{h : [N] \rightarrow [M]\}$ is *k-wise independent* if for all distinct $x_1, x_2, \dots, x_k \in [N]$, the random variables $H(x_1), \dots, H(x_k)$ are independent and uniformly distributed in $[M]$ when $H \stackrel{R}{\leftarrow} \mathcal{H}$.

Construction 3.32 (*k-wise independence from polynomials*). Let \mathbb{F} be a finite field. Define the family of functions $\mathcal{H} = \{h_{a_0, a_1, \dots, a_{k-1}} : \mathbb{F} \rightarrow \mathbb{F}\}$ where each $h_{a_0, a_1, \dots, a_{k-1}}(x) = a_0 + a_1x + a_2x^2 + \dots + a_{k-1}x^{k-1}$ for $a_0, \dots, a_{k-1} \in \mathbb{F}$.

Proposition 3.33. The family \mathcal{H} given in Construction 3.32 is *k-wise independent*.

Proof. Similarly to the proof of Proposition 3.24, it suffices to prove that for all distinct $x_1, \dots, x_k \in \mathbb{F}$ and all $y_1, \dots, y_k \in \mathbb{F}$, there is exactly one polynomial h of degree at most $k-1$ such that $h(x_i) = y_i$ for all i . To show that such a polynomial exists, we can use the Lagrange Interpolation Formula:

$$h(x) = \sum_{i=1}^k y_i \cdot \prod_{j \neq i} \frac{x - x_j}{x_i - x_j}.$$

To show uniqueness, suppose we have two polynomials h and g of degree at most $k-1$ such that $h(x_i) = g(x_i)$ for $i = 1, \dots, k$. Then $h-g$ has at least k roots, and thus must be the zero polynomial. \square

Corollary 3.34. For every $n, m, k \in \mathbb{N}$, there is a family of *k-wise independent* functions $\mathcal{H} = \{h : \{0, 1\}^n \rightarrow \{0, 1\}^m\}$ such that choosing a random function from \mathcal{H} takes $k \cdot \max\{n, m\}$ random bits, and evaluating a function from \mathcal{H} takes time $\text{poly}(n, m, k)$.

k-wise independent hash functions have applications similar to those that pairwise independent hash functions have. The increased independence is crucial in derandomizing some algorithms. *k-wise independent* random variables also satisfy a tail bound similar to Proposition 3.28, with the key improvement being that the error probability vanishes linearly in $t^{k/2}$ rather than t ; see Problem 3.8.

3.6 Exercises

Problem 3.1 (Derandomizing RP versus BPP*). Show that $\text{prRP} = \text{prP}$ implies that $\text{prBPP} = \text{prP}$, and thus also that $\text{BPP} = \text{P}$. (Hint: Look at the proof that $\text{NP} = \text{P} \Rightarrow \text{BPP} = \text{P}$.)

Problem 3.2 (Designs). Designs (also known as packings) are collections of sets that are nearly disjoint. In Chapter 7, we will see how they are useful in the construction of pseudorandom generators. Formally, a collection of sets $S_1, S_2, \dots, S_m \subset [d]$ is called an (ℓ, a) -design (for integers $a \leq \ell \leq d$) if

- For all i , $|S_i| = \ell$.
- For all $i \neq j$, $|S_i \cap S_j| < a$.

For given ℓ , we'd like m to be large, a to be small, and d to be small. That is, we'd like to pack many sets into a small universe with small intersections.

- (1) Prove that if $m \leq \binom{d}{a} / \binom{\ell}{a}^2$, then there exists an (ℓ, a) -design $S_1, \dots, S_m \subset [d]$.

Hint: Use the Probabilistic Method. Specifically, show that if the sets are chosen randomly, then for every S_1, \dots, S_{i-1} ,

$$\mathbb{E}_{S_i} [\#\{j < i : |S_i \cap S_j| \geq a\}] < 1.$$

- (2) Conclude that for every constant $\gamma > 0$ and every $\ell, m \in \mathbb{N}$, there exists an (ℓ, a) -design $S_1, \dots, S_m \subset [d]$ with $d = O\left(\frac{\ell^2}{a}\right)$ and $a = \gamma \cdot \log m$. In particular, setting $m = 2^\ell$, we fit exponentially many sets of size ℓ in a universe of size $d = O(\ell)$ while keeping the intersections an arbitrarily small fraction of the set size.
- (3) Using the Method of Conditional Expectations, show how to construct designs as in Parts 1 and 2 *deterministically* in time $\text{poly}(m, d)$.
-

Problem 3.3 (More Pairwise Independent Families). (1) (matrix-vector family) For an $n \times m$ $\{0, 1\}$ -matrix A and $b \in \{0, 1\}^n$, define a function $h_{A,b} : \{0, 1\}^m \rightarrow \{0, 1\}^n$ by $h_{A,b}(x) = (Ax + b) \bmod 2$. (The “mod 2” is applied componentwise.) Show that $\mathcal{H}_{m,n} = \{h_{A,b}\}$ is a pairwise independent family. Compare the number of random bits needed to generate a random function in $\mathcal{H}_{m,n}$ to Construction 3.23.

- (2) (Toeplitz matrices) A is a *Toeplitz matrix* if it is constant on diagonals, i.e. $A_{i+1,j+1} = A_{i,j}$ for all i, j . Show that even if we restrict the family $\mathcal{H}_{m,n}$ in Part 1 to only include $h_{A,b}$ for Toeplitz matrices A , we still get a pairwise independent family. How many random bits are needed now?
-

Problem 3.4 (Almost Pairwise Independence). A family of functions \mathcal{H} mapping domain $[N]$ to range $[M]$ is ε -almost pairwise independent³ if for every $x_1 \neq x_2 \in [N]$, $y_1, y_2 \in [M]$, we have

$$\Pr_{H \leftarrow \mathcal{H}} [H(x_1) = y_1 \text{ and } H(x_2) = y_2] \leq \frac{1 + \varepsilon}{M^2}.$$

³Another common definition of ε -almost pairwise independence requires instead that for every $x_1 \neq x_2 \in [N]$, if we choose a random hash function $H \leftarrow \mathcal{H}$, the random variable $(H(x_1), H(x_2))$ is ε -close to two uniform and independent elements of $[M]$ in statistical difference (as defined in Chapter 6). The two definitions are equivalent up to a factor of M^2 in the error parameter ε .

- (1) Show that there exists a family \mathcal{H} of ε -almost pairwise independent functions from $\{0, 1\}^n$ to $\{0, 1\}^m$ such that choosing a random function from \mathcal{H} requires only $O(m + \log n + \log(1/\varepsilon))$ random bits (as opposed to $O(m + n)$ for exact pairwise independence). (Hint: First consider domain \mathbb{F}^{d+1} for an appropriately chosen finite field \mathbb{F} and $d \in \mathbb{N}$, and look at maps of the form $h = g \circ f_a$, where g comes from some pairwise independent family and $f_a : \mathbb{F}^{d+1} \rightarrow \mathbb{F}$ is defined by $f_a(x_0, \dots, x_d) = x_0 + x_1 a + x_2 a^2 + \dots + x_d a^d$.)
- (2) Give a deterministic algorithm that on input an N -vertex, M -edge graph G (with no self-loops), finds a cut of size at least $(1/2 - o(1)) \cdot M$ in time $M \cdot \text{polylog}(N)$ and space $O(\log M)$ (thereby improving the $M \cdot \text{poly}(N)$ running time of Algorithm 3.20).

Problem 3.5 (Size Lower Bound for Pairwise Independent Families). Let $\mathcal{H} = \{h : [N] \rightarrow [M]\}$ be a pairwise independent family of functions.

- (1) Prove that if $N \geq 2$, then $|\mathcal{H}| \geq M^2$.
- (2) Prove that if $M = 2$, then $|\mathcal{H}| \geq N + 1$. (Hint: based on \mathcal{H} , construct a sequence of orthogonal vectors $v_x \in \{\pm 1\}^{|\mathcal{H}|}$ parameterized by $x \in [N]$.)
- (3) More generally, prove that for arbitrary M , we have $|\mathcal{H}| \geq N \cdot (M - 1) + 1$. (Hint: for each $x \in [N]$, construct $M - 1$ linearly independent vectors $v_{x,y} \in \mathbb{R}^{|\mathcal{H}|}$ such that $v_{x,y} \perp v_{x',y'}$ if $x \neq x'$.)
- (4) Deduce that for $N = 2^n$ and $M = 2^m$, selecting a random function from \mathcal{H} requires at least $\max\{n, m\} + m$ random bits.

Problem 3.6 (Frequency Moments of Data Streams). Given one pass through a huge “stream” of data items (a_1, a_2, \dots, a_k) , where each $a_i \in \{0, 1\}^n$, we want to compute statistics on the distribution of items occurring in the stream while using small space (not enough to store all the items or maintain a histogram). In this problem, you will see how to compute the *2nd frequency moment* $f_2 = \sum_a m_a^2$, where $m_a = \#\{i : a_i = a\}$.

The algorithm works as follows: Before receiving any items, it chooses t random *4-wise independent* hash functions $H_1, \dots, H_t : \{0, 1\}^n \rightarrow \{+1, -1\}$, and sets counters $X_1 = X_2 = \dots = X_t = 0$. Upon receiving the i 'th item a_i , it adds $H_j(a_i)$ to counter X_j . At the end of the stream, it outputs $Y = (X_1^2 + \dots + X_t^2)/t$.

Notice that the algorithm only needs space $O(t \cdot n)$ to store the hash functions H_j and space $O(t \cdot \log k)$ to maintain the counters X_j (compared to space $k \cdot n$ to store the entire stream, and space $2^n \cdot \log k$ to maintain a histogram).

- (1) Show that for every data stream (a_1, \dots, a_k) and each j , we have $\mathbb{E}[X_j^2] = f_2$, where the expectation is over the choice of the hash function H_j .
- (2) Show that $\text{Var}[X_j^2] \leq 2f_2^2$.
- (3) Conclude that for a sufficiently large constant t (independent of n and k), the output Y is within 1% of f_2 with probability at least .99.

- (4) Show how to decrease the error probability to δ while only increasing the space by a factor of $\log(1/\delta)$.
-

Problem 3.7 (Improved Pairwise Independent Tail Inequality). (1) Show that if X is a random variable taking values in $[0, 1]$ and $\mu = \mathbb{E}[X]$, we have $\text{Var}[X] \leq \mu \cdot (1 - \mu)$.
 (2) Improve the Pairwise Independent Tail Inequality (Proposition 3.28) to show that if X is the average of t pairwise independent random variables taking values in $[0, 1]$ and $\mu = \mathbb{E}[X]$, then $\Pr[|X - \mu| \geq \varepsilon] \leq \mu \cdot (1 - \mu)/(t \cdot \varepsilon^2)$. In particular, for $t = O(1/\mu)$, we have $\Pr[.99\mu \leq X \leq 1.01\mu] \geq .99$.

Problem 3.8 (k -wise Independent Tail Inequality). Let X be the average of t k -wise independent random variables for an even integer k , and let $\mu = \mathbb{E}[X]$. Prove that

$$\Pr[|X - \mu| \geq \varepsilon] \leq \left(\frac{k^2}{4t\varepsilon^2} \right)^{k/2}.$$

(Hint: show that all terms in the expansion of $\mathbb{E}[(X - \mu)^k] = \mathbb{E}[(\sum_i (X_i - \mu_i))^k]$ that involve more than $k/2$ variables X_i are zero.) Note that for fixed k , this probability decays like $O(1/(t\varepsilon^2))^{k/2}$, improving the $1/(t\varepsilon^2)$ bound in pairwise independence when $k > 2$.

Problem 3.9 (Hitting Samplers). A function $\text{Samp} : [N] \rightarrow [M]^t$ is a (δ, ε) *hitting sampler* if for every set $T \subset [M]$ of density greater than ε , we have

$$\Pr_{(z_1, \dots, z_t) \stackrel{\mathbb{R}}{\leftarrow} \text{Samp}(U_{[N]})} [\exists i \quad z_i \in T] \geq 1 - \delta.$$

- (1) Show that every (δ, ε) averaging sampler is a (δ, ε) hitting sampler.
 - (2) Show that if we only want a hitting sampler, the number of samples in Theorem 3.30 can be reduced to $O(1/(\varepsilon\delta))$. (Hint: use Problem 3.7.)
 - (3) For which subset of **BPP** algorithms are hitting samplers useful for doing randomness-efficient error reduction?
-

3.7 Chapter Notes and References

The Time Hierarchy Theorem was proven by Hartmanis and Stearns [HS]; proofs can be found in any standard text on complexity theory, e.g. [Sip3, Gol7, AB]. Adleman [Adl] showed that every language in **RP** has polynomial-sized circuits (cf., Corollary 3.12), and this was generalized to **BPP** by Gill. Pippenger [Pip] showed the equivalence between having polynomial-sized circuits and **P/poly** (Fact 3.11). The general definition of complexity classes with advice (Definition 3.10) is due to Karp and Lipton [KL], who explored the relationship between nonuniform lower bounds

and uniform lower bounds. A $5n - o(n)$ circuit-size lower bound for an explicit function (in \mathbf{P}) was given by Iwama et al. [LR, IM].

The existence of universal traversal sequences (Example 3.8) was proven by Aleliunas et al. [AKL⁺], who suggested finding an explicit construction (Open Problem 3.9) as an approach to derandomizing the logspace algorithm for `UNDIRECTED S-T CONNECTIVITY`. For the state of the art on these problems, see Section 4.4. An conjectured deterministic \mathbf{NC} algorithm for `PERFECT MATCHING` (derandomizing Algorithm 2.7 in a different way than Open Problem 3.7) is given in [Agr].

Theorem 3.14 is due to Sipser [Sip1], who proved that \mathbf{BPP} is in the 4th level of the polynomial-time hierarchy; this was improved to the 2nd level by Gács. Our proof of Theorem 3.14 is due to Lautemann [Lau]. Problem 3.1 is due to Buhrman and Fortnow [BF]. For more on nondeterministic computation and nonuniform complexity, see textbooks on computational complexity, such as [Sip3, Gol7, AB].

The Method of Conditional Probabilities was formalized and popularized as an algorithmic tool in the work of Spencer [Spe] and Raghavan [Rag]. Its use in Algorithm 3.17 for approximating `MAXCUT` is implicit in [Lub2]. For more on this method, see the textbooks [MR, AS1].

A more detailed treatment of pairwise independence (along with a variety of other topics in pseudorandomness and derandomization) can be found in the survey by Luby and Wigderson [LW]. The use of limited independence in computer science originates with the seminal papers of Carter and Wegman [CW, WC], which introduced the notions of universal, strongly universal (i.e. k -wise independent), and almost strongly universal (i.e. almost k -wise independent) families of hash functions. The pairwise independent and k -wise independent sample spaces of Constructions 3.18, 3.23, and 3.32 date back to the work of Lancaster [Lan] and Joffe [Jof1, Jof2] in the probability literature, and were rediscovered several times in the computer science literature. The construction of pairwise independent hash functions from Part 1 of Problem 3.3 is due to Carter and Wegman [CW] and Part 2 is implicit in [Vaz2, GL]. The size lower bound for pairwise independent families in Problem 3.5 is due to Stinson [Sti2], based on the Plackett–Burman bound for orthogonal arrays [PB]. The construction of almost pairwise independent families in Problem 3.4 is due to Bierbrauer et al. [BJKS] (though the resulting parameters follow from the earlier work of Naor and Naor [NN]).

The application to hash tables from Section 3.5.3 is due to Carter and Wegman [CW], and the method mentioned for improving the space complexity to $O(K \log N)$ is due to Fredman, Komlós, and Szemerédi [FKS]. The problem of randomness-efficient error reduction (sometimes called “deterministic amplification”) was first studied by Karp, Pippenger, and Sipser [KPS], and the method using pairwise independence given in Section 3.5.4 was proposed by Chor and Goldreich [CG2]. The use of pairwise independence for derandomizing algorithms was pioneered by Luby [Lub1]; Algorithm 3.20 for `MAXCUT` is implicit in [Lub2]. The notion of averaging samplers was introduced by Bellare and Rompel [BR] (under the term “oblivious samplers”). For more on samplers and averaging samplers, see the survey by Goldreich [Gol1]. Tail bounds for k -wise independent random variables, such as the one in Problem 3.8, can be found in the papers [CG2, BR, SSS].

Problem 3.2 on designs is from [EFF], with the derandomization of Part 3 being from [NW, LW]. Problem 3.6 on the frequency moments of data streams is due to Alon, Matias, and Szegedy [AMS]. For more on data stream algorithms, we refer to the survey by Muthukrishnan [Mut].

4

Expander Graphs

Now that we have seen a variety of basic derandomization techniques, we will move on to study the first major “pseudorandom object” in this survey, *expander graphs*. These are graphs that are “sparse” yet very “well-connected.”

4.1 Measures of Expansion

We will typically interpret the properties of expander graphs in an asymptotic sense. That is, there will be an infinite family of graphs G_i , with a growing number of vertices N_i . By “sparse,” we mean that the degree D_i of G_i should be very slowly growing as a function of N_i . The “well-connectedness” property has a variety of different interpretations, which we will discuss below. Typically, we will drop the subscripts of i and the fact that we are talking about an infinite family of graphs will be implicit in our theorems. As in Section 2.4.2, we will state many of our definitions for *directed multigraphs* (which we’ll call *digraphs* for short), though in the end we will mostly study undirected multigraphs.

4.1.1 Vertex Expansion

The classic measure of well-connectedness in expanders requires that every “not-too-large” set of vertices has “many” neighbors:

Definition 4.1. A digraph G is a (K, A) *vertex expander* if for all sets S of at most K vertices, the *neighborhood* $N(S) \stackrel{\text{def}}{=} \{u | \exists v \in S \text{ s.t. } (u, v) \in E\}$ is of size at least $A \cdot |S|$.

Ideally, we would like $D = O(1)$, $K = \Omega(N)$ where N is the number of vertices, and A as close to D as possible.

There are several other measures of expansion, some of which we will examine in forthcoming sections:

- **Boundary expansion:** instead of $N(S)$, only use the *boundary* $\partial S \stackrel{\text{def}}{=} N(S) \setminus S$.

- Edge expansion (cuts): instead of ∂S , use the number of edges leaving S .
- Random walks: random walks converge quickly to the uniform distribution, i.e. the second eigenvalue $\lambda(G)$ is small.
- “Quasi-randomness” (a.k.a “Mixing”): for every two sets S and T (say of constant density), the fraction of edges between S and T is roughly the product of their densities.

All of these measures are very closely related to each other, and are even equivalent for certain settings of parameters.

It is not obvious from the definition that good vertex expanders (say, with $D = O(1)$, $K = \Omega(N)$, and $A = 1 + \Omega(1)$) even exist. We will show this using the Probabilistic Method.

Theorem 4.2. For all constants $D \geq 3$, there is a constant $\alpha > 0$ such that for all N , a random D -regular undirected graph on N vertices is an $(\alpha N, D - 1.01)$ vertex expander with probability at least $1/2$.

Note that the degree bound of 3 is the smallest possible, as every graph of degree 2 is a poor expander (being a union of cycles and chains). In addition, for most settings of parameters, it is impossible to have expansion larger than $D - 1$ (as shown in Problem 4.3).

We prove a slightly simpler theorem for *bipartite expanders*.

Definition 4.3. A bipartite multigraph G is a (K, A) *vertex expander* if for all sets S of *left*-vertices of size at most K , the neighborhood $N(S)$ is of size at least $A \cdot |S|$.

Now, let $\text{Bip}_{N,D}$ be the set of bipartite multigraphs that have N vertices on each side and are *D-leftregular*, meaning that every vertex on the left has D neighbors, numbered from $1, \dots, D$ (but vertices on the right may have varying degrees).

Theorem 4.4. For every constant D , there exists a constant $\alpha > 0$ such that for all N , a uniformly random graph from $\text{Bip}_{N,D}$ is an $(\alpha N, D - 2)$ vertex expander with probability at least $1/2$.

Proof. First, note that choosing $G \stackrel{\text{R}}{\leftarrow} \text{Bip}_{N,D}$ is equivalent to uniformly and independently choosing D neighbors on the right for each left vertex v . Now, for $K \leq \alpha N$, let p_K be the probability that there exists a left-set S of size exactly K that does not expand by at least $D - 2$. Fixing a subset S of size K , $N(S)$ is a set of KD random vertices in $[N]$ (chosen with replacement). We can imagine these vertices V_1, V_2, \dots, V_{KD} being chosen in sequence. Call V_i a *repeat* if $V_i \in \{V_1, \dots, V_{i-1}\}$. Then the probability that V_i is a repeat, even conditioned on V_1, \dots, V_{i-1} , is at most $(i-1)/N \leq KD/N$. So,

$$\begin{aligned} \Pr[|N(S)| \leq (D - 2) \cdot K] &\leq \Pr[\text{there are at least } 2K \text{ repeats among } V_1, \dots, V_{KD}] \\ &\leq \binom{KD}{2K} \left(\frac{KD}{N}\right)^{2K}. \end{aligned}$$

Thus, we find that

$$p_K \leq \binom{N}{K} \binom{KD}{2K} \left(\frac{KD}{N}\right)^{2K} \leq \left(\frac{Ne}{K}\right)^K \left(\frac{KDe}{2K}\right)^{2K} \left(\frac{KD}{N}\right)^{2K} = \left(\frac{e^3 D^4 K}{4N}\right)^K$$

where e is the base of the natural logarithm. Since $K \leq \alpha N$, we can set $\alpha = 1/(e^3 D^4)$ to obtain $p_K \leq 4^{-K}$. Thus

$$\Pr_{G \in \text{Bip}_{N,D}} [G \text{ is not an } (\alpha N, D-2) \text{ expander}] \leq \sum_{K=1}^{\lfloor \alpha N \rfloor} 4^{-K} < \frac{1}{2}$$

□

There are a number of variants to the above probabilistic construction of expanders.

- We can obtain a bipartite multigraph that is D -regular on both sides by taking the union of D random perfect matchings. This can be analyzed using a small modification of the analysis above; even though V_1, \dots, V_{KD} are not independent, the probability of a V_i being a repeat conditioned on V_1, \dots, V_{i-1} can still be bounded by $KD/(N-K)$. Also, the multigraph can be made into a simple graph by eliminating or redistributing edges.
- One can optimize α rather than the expansion factor A , showing that for all constants $\alpha < 1$ and $D > 2$, there exists a constant $A > 1$ such that for all sufficiently large N , a random graph in $\text{Bip}_{N,D}$ is an $(\alpha N, A)$ vertex expander with high probability.
- In fact, a very general tradeoff between D , α , and A is known: a random D -regular N -vertex bipartite multigraph is an $(\alpha N, A)$ vertex expander with high probability for sufficiently large N if $D > \frac{H(\alpha) + H(\alpha A)}{H(\alpha) - \alpha A H(1/A)}$, where $H(p) = p \log(1/p) + (1-p) \log(1/(1-p))$ is the binary entropy function.
- The results can also be extended to unbalanced bipartite graphs (where the right side is smaller than the left), and nonbipartite graphs as well, and both of these cases are important in some applications.

In addition to being natural combinatorial objects, expander graphs have numerous applications in theoretical computer science, including the construction of fault-tolerant networks (indeed, the first papers on expanders were in conferences on telephone networks), sorting in $O(\log n)$ time in parallel, derandomization (as we will see), lower bounds in circuit complexity and proof complexity, error-correcting codes, negative results regarding integrality gaps for linear programming relaxations and metric embeddings, distributed routing, and data structures. For many of these applications, it is not enough to know that a random graph is a good expander — we need *explicit* constructions, i.e. ones that are deterministic and efficient. We view such explicit expanders as “pseudorandom objects” because they are fixed graphs that possess many of the properties of random graphs.

4.1.2 Spectral Expansion

Intuitively, another way of saying that a graph is well-connected is to require that random walks on the graph converge quickly to the stationary distribution. As we have seen in Section 2.4.2, the mixing rate of random walks in turn is captured well by the second largest eigenvalue of the transition matrix, and this turns out to be a very useful measure of expansion for many purposes.

Recall that for an N -vertex regular directed graph G with random-walk matrix M , we define

$$\lambda(G) \stackrel{\text{def}}{=} \max_{\pi} \frac{\|\pi M - \pi\|}{\|\pi - \pi\|} = \max_{x \perp u} \frac{\|xM\|}{\|x\|},$$

where $u = (1/N, \dots, 1/N) \in \mathbb{R}^N$ is the uniform distribution on $[N]$, the first maximum is over all probability distributions $\pi \in [0, 1]^N$, and the second maximum is over all vectors $x \in \mathbb{R}^N$ that are orthogonal to u . We write $\gamma(G) \stackrel{\text{def}}{=} 1 - \lambda(G)$ to denote the *spectral gap* of G .

Definition 4.5. For $\gamma \in [0, 1]$, we say that a regular digraph G has *spectral expansion* γ if $\gamma(G) \geq \gamma$ (equivalently, $\lambda(G) \leq 1 - \gamma$).¹

Larger values of $\gamma(G)$ (or smaller values of $\lambda(G)$) correspond to better expansion. Sometimes it is more natural to state results in terms of $\gamma(G)$ and sometimes in terms of $\lambda(G)$. Surprisingly, this linear-algebraic measure of expansion turns out to be equivalent to the combinatorial measure of vertex expansion for common parameters of interest.

One direction is given by the following:

Theorem 4.6 (spectral expansion \Rightarrow vertex expansion). If G is a regular digraph with spectral expansion $\gamma = 1 - \lambda$ for some $\lambda \in [0, 1]$, then, for every $\alpha \in [0, 1]$, G is an $(\alpha N, 1/((1 - \alpha)\lambda^2 + \alpha))$ vertex expander. In particular, G is an $(N/2, 1 + \gamma)$ expander.

We prove this theorem using the following two useful statistics of probability distributions.

Definition 4.7. For a probability distribution π , the *collision probability* of π is defined to be the probability that two independent samples from π are equal, namely $\text{CP}(\pi) = \sum_x \pi_x^2$. The *support* of π is $\text{Supp}(\pi) = \{x : \pi_x > 0\}$.

Lemma 4.8. For every probability distribution $\pi \in [0, 1]^N$, we have:

- (1) $\text{CP}(\pi) = \|\pi\|^2 = \|\pi - u\|^2 + 1/N$, where u is the uniform distribution on $[N]$.
 - (2) $\text{CP}(\pi) \geq 1/|\text{Supp}(\pi)|$, with equality iff π is uniform on $\text{Supp}(\pi)$.
-

Proof. For Part 1, the fact that $\text{CP}(\pi) = \|\pi\|^2$ follows immediately from the definition. Then, writing $\pi = u + (\pi - u)$, and noting that $\pi - u$ is orthogonal to u , we have $\|\pi\|^2 = \|u\|^2 + \|\pi - u\|^2 = 1/N + \|\pi - u\|^2$.

For Part 2, by Cauchy–Schwarz we have

$$1 = \sum_{x \in \text{Supp}(\pi)} \pi_x \leq \sqrt{|\text{Supp}(\pi)|} \cdot \sqrt{\sum_x \pi_x^2} = \sqrt{|\text{Supp}(\pi)|} \cdot \sqrt{\text{CP}(\pi)},$$

with equality iff π is uniform on $\text{Supp}(\pi)$. □

¹In other sources (including the original lecture notes on which this survey was based), the spectral expansion referred to λ rather than γ . Here we use γ , because it has the more natural feature that larger values of γ correspond to the graph being “more expanding”.

Proof of Theorem 4.6. The condition that G has spectral expansion $\gamma = 1 - \lambda$ is equivalent to saying that $\lambda(G) \leq \lambda$. By the definition of $\lambda(G)$ and Part 1 of Lemma 4.8, we have

$$\text{CP}(\pi M) - \frac{1}{N} \leq \lambda^2 \cdot \left(\text{CP}(\pi) - \frac{1}{N} \right)$$

for every probability distribution π . Letting S be any subset of the vertices of size at most αN and π the uniform distribution on S , we have $\text{CP}(\pi) = 1/|S|$ and $\text{CP}(\pi M) \geq 1/|\text{Supp}(\pi M)| = 1/|N(S)|$. Thus,

$$\left(\frac{1}{|N(S)|} - \frac{1}{N} \right) \leq \lambda^2 \cdot \left(\frac{1}{|S|} - \frac{1}{N} \right)$$

Solving for $|N(S)|$ and using $N \geq |S|/\alpha$, we obtain $|N(S)| \geq |S|/(\lambda^2(1 - \alpha) + \alpha)$, as desired. \square

The other direction, i.e. obtaining spectral expansion from vertex expansion, is more difficult (and we will not prove it here).

Theorem 4.9 (vertex expansion \Rightarrow spectral expansion). For every $\delta > 0$ and $D > 0$, there exists $\gamma > 0$ such that if G is a D -regular $(N/2, 1 + \delta)$ vertex expander, then it also has spectral expansion γ . Specifically, we can take $\gamma = \Omega((\delta/D)^2)$.

Note first the dependence on subset size being $N/2$: this is necessary, because a graph can have vertex expansion $(\alpha N, 1 + \Omega(1))$ for $\alpha < 1/2$ and be disconnected (e.g. the disjoint union of two good expanders), thereby having no spectral expansion. Also note that the bound on γ deteriorates as D increases. This is also necessary, because adding edges to a good expander cannot hurt its vertex expansion, but can hurt its spectral expansion.

Still, roughly speaking, these two results show that vertex expansion and spectral expansion are closely related, indeed equivalent for many interesting settings of parameters:

Corollary 4.10. Let \mathcal{G} be an infinite family of D -regular multigraphs, for a constant $D \in \mathbb{N}$. Then the following two conditions are equivalent:

- There is a constant $\delta > 0$ such that every $G \in \mathcal{G}$ is an $(N/2, 1 + \delta)$ vertex expander.
- There is a constant $\gamma > 0$ such that every $G \in \mathcal{G}$ has spectral expansion γ .

When people informally use the term “expander,” they often mean a family of regular graphs of *constant degree* D satisfying one of the two equivalent conditions above.

However, the two measures are no longer equivalent if one wants to optimize the expansion constants. For vertex expansion, we have already seen that if we allow α to be a small constant (depending on D), then there exist $(\alpha N, A)$ vertex expanders with A very close to $D - 1$, e.g. $A = D - 1.01$, and Problem 4.3 shows that A cannot be any larger than $D - 1$. The optimal value for the spectral expansion is also well-understood. First note that, by taking $\alpha \rightarrow 0$ in Theorem 4.6, a graph with spectral expansion $1 - \lambda$ has vertex expansion $A \approx 1/\lambda^2$ for small sets. Thus, a lower bound on λ is $1/\sqrt{D} - o(1)$. In fact, this lower bound can be improved, as shown in the following theorem (and proven in Problem 4.4):

Theorem 4.11. For every constant $D \in \mathbb{N}$, any D -regular, N -vertex multigraph G satisfies $\lambda(G) \geq 2\sqrt{D-1}/D - o(1)$, where the $o(1)$ term vanishes as $N \rightarrow \infty$ (and D is held constant).

Surprisingly, there exist explicit constructions giving $\lambda(G) < 2\sqrt{D-1}/D$. Graphs meeting this bound are called *Ramanujan graphs*. Random graphs almost match this bound, as well:

Theorem 4.12. For any constant $D \in \mathbb{N}$, a random D -regular N -vertex graph satisfies $\lambda(G) \leq 2\sqrt{D-1}/D + o(1)$ with probability $1 - o(1)$ where both $o(1)$ terms vanish as $N \rightarrow \infty$ (and D is held constant).

Now let us see what these results for spectral expansion imply in the world of vertex expansion. With Ramanujan graphs ($\lambda(G) \leq 2\sqrt{D-1}/D$), the bound from Theorem 4.6 gives a vertex expansion factor of $A \approx D/4$ for small sets. This is not tight, and it is known that Ramanujan graphs actually have vertex expansion $D/2 - o(1)$ for sets of density $o(1)$, which is tight in the sense that there are families of graphs with $\lambda(G) \rightarrow 2\sqrt{D-1}/D$ but vertex expansion at most $D/2$. Still, this vertex expansion is not as good as we obtained via the Probabilistic Method (Theorem 4.2), where we achieved vertex expansion $D - O(1)$. This means that we cannot obtain optimal vertex expansion by going through spectral expansion. Similarly, we cannot obtain optimal spectral expansion by going through vertex expansion (because the bound on spectral expansion in Theorem 4.9 necessarily deteriorates as the degree D increases). The conclusion is that vertex and spectral expansion are loosely equivalent, but only if we are not interested in optimizing the constants in the tradeoffs between various parameters (and for some applications these are crucial).

4.1.3 Other Measures of Expansion

In this section, we mention two other useful measures of expansion involving edges crossing cuts in the graph. For two sets $S, T \subset V(G)$, let $e(S, T) = \|\{(u, v) \in S \times T \mid \{u, v\} \in E\}\|$. Here (u, v) refers to an *ordered* pair, in contrast to the definition of $\text{cut}(S, T)$ in Section 2.3.4. Thus, we count edges entirely within $S \cap T$ twice, corresponding to both orientations.

Definition 4.13. A D -regular digraph G is a (K, ε) *edge expander* if for all sets S of at most K vertices, the cut size $e(S, \bar{S})$ is at least $\varepsilon \cdot |S| \cdot D$.

That is, at least an ε fraction of the edges from S lead outside S . (Sometimes edge expansion is defined without the normalization factor of D , only requiring $e(S, \bar{S}) \geq \varepsilon \cdot |S|$.) When viewed in terms of the random walk on G , the ratio $e(S, \bar{S})/(|S| \cdot D)$ is the probability that, if we condition the stationary distribution on being in S , the random walk leaves S in one step. It turns out that if we fix $K = N/2$, then edge expansion turns out to be even more closely related to spectral expansion than vertex expansion is. Indeed:

Theorem 4.14.

- (1) If a D -regular, N -vertex digraph G has spectral expansion γ , then G is an $(N/2, \gamma/2)$ edge expander.
- (2) If a D -regular, N -vertex digraph G is a $(N/2, \varepsilon)$ edge expander and at least an α fraction of edges leaving each vertex are self-loops for some $\alpha \in [0, 1]$, then G has spectral expansion $\alpha \cdot \varepsilon^2/2$.

The condition about self-loops in Part 2 is to ensure that the graph is far from being bipartite (or more generally “periodic” in the sense that all cycle lengths are divisible by some number larger than 1), because a bipartite graph has spectral expansion 0 but can have positive edge expansion. For graphs with a constant fraction of self-loops at each vertex, the theorem implies that the edge expansion is bounded away from 0 iff the spectral expansion is bounded away from 0. Unlike Corollary 4.10, this equivalence holds even for graphs of unbounded degree. The intuition for the relation is that a large edge expansion ε implies that the random walk on the graph has no “bottlenecks” and thus should mix rapidly. This connection also holds for Markov chains in general (when the definitions are appropriately generalized), where the edge expansion is known as the *conductance*. Part 1 of Theorem 4.14 will follow as a special case of the Expander Mixing Lemma below; we omit the proof of Part 2.

Next, we consider a generalization of edge expansion, where we look at edges not just from a set S to its complement but between any two sets S and T . If we think of an expander as being like a random graph, we would expect the fraction of edges that go from S to T to be approximately equal to the product of the densities of S and T . The following result shows that this intuition is correct:

Lemma 4.15 (Expander Mixing Lemma). Let G be a D -regular, N -vertex digraph with spectral expansion $1 - \lambda$. Then for all sets of vertices S, T of densities $\alpha = |S|/N$ and $\beta = |T|/N$, we have

$$\begin{aligned} \left| \frac{e(S, T)}{N \cdot D} - \alpha\beta \right| &\leq \lambda \sqrt{\alpha \cdot (1 - \alpha) \cdot \beta \cdot (1 - \beta)}. \\ &\leq \lambda \sqrt{\alpha\beta} \leq \lambda. \end{aligned}$$

Observe that the denominator $N \cdot D$ counts all edges of the graph (as ordered pairs). The lemma states that the difference between the fraction of edges from S to T and the expected value if we were to choose G randomly is “small”, roughly λ times the square root of this fraction. Finally, note that Part 1 of Theorem 4.14 follows from the Expander Mixing Lemma by setting $T = S^c$, so $\beta = 1 - \alpha$ and $e(S, T)/ND \geq (1 - \lambda) \cdot \alpha \cdot (1 - \alpha) \geq \gamma\alpha/2$.

When a digraph $G = (V, E)$ has the property that $|e(S, T)/|E| - \alpha\beta| = o(1)$ for all sets S, T (with densities α, β), the graph is called *quasirandom*. Thus, the Expander Mixing Lemma implies that a regular digraph with $\lambda(G) = o(1)$ is quasirandom. Quasirandomness has been studied extensively for dense graphs, in which case it has numerous equivalent formulations. Here we are most interested in sparse graphs, especially constant-degree graphs (for which $\lambda(G) = o(1)$ is impossible).

Proof. Let χ_S be the characteristic (row) vector of S and χ_T the characteristic vector of T . Let A be the adjacency matrix of G , and $M = A/D$ be the random-walk matrix for G . Note that $e(S, T) = \chi_S A \chi_T^t = \chi_S (DM) \chi_T^t$, where the superscript t denotes the transpose.

We can express χ_S as the sum of two components, one parallel to the uniform distribution u , and the other a vector χ_S^\perp , where $\chi_S^\perp \perp u$. The coefficient of u is $\langle \chi_S, u \rangle / \langle u, u \rangle = \sum_i (\chi_S)_i = |S| = \alpha N$. Then $\chi_S = (\alpha N)u + \chi_S^\perp$ and similarly $\chi_T = (\beta N)u + \chi_T^\perp$. Intuitively, the components parallel to the uniform distribution “spread” the weight of S and T uniformly over the entire graph, and χ_S^\perp and χ_T^\perp will yield the error term.

Formally, we have

$$\begin{aligned} \frac{e(S, T)}{N \cdot D} &= \frac{1}{N} ((\alpha N)u + \chi_S^\perp) M ((\beta N)u + \chi_T^\perp)^t \\ &= \frac{1}{N} (\alpha \beta N^2) u M u^t + \frac{1}{N} (\alpha N) u M (\chi_T^\perp)^t + \frac{1}{N} (\beta N) \chi_S^\perp M u^t + \frac{1}{N} \chi_S^\perp M (\chi_T^\perp)^t. \end{aligned}$$

Since $uM = u$ and $Mu^t = u^t$, and both χ_S^\perp and χ_T^\perp are orthogonal to u , the above expression simplifies to:

$$\frac{e(S, T)}{N \cdot D} = (\alpha \beta N) u u^t + \frac{\chi_S^\perp M (\chi_T^\perp)^t}{N} = \alpha \beta + \frac{(\chi_S^\perp M) (\chi_T^\perp)^t}{N}.$$

Thus,

$$\begin{aligned} \left| \frac{e(S, T)}{N \cdot D} - \alpha \beta \right| &= \left| \frac{(\chi_S^\perp M) (\chi_T^\perp)^t}{N} \right| \\ &\leq \frac{1}{N} \cdot \|\chi_S^\perp M\| \cdot \|\chi_T^\perp\| \\ &\leq \frac{1}{N} \cdot \lambda \|\chi_S^\perp\| \cdot \|\chi_T^\perp\|. \end{aligned}$$

To complete the proof, we note that

$$\alpha N = \|\chi_S\|^2 = \|(\alpha N)u\|^2 + \|\chi_S^\perp\|^2 = \alpha^2 N + \|\chi_S^\perp\|^2,$$

so $\|\chi_S^\perp\| = \sqrt{(\alpha - \alpha^2)N} = \sqrt{\alpha \cdot (1 - \alpha) \cdot N}$ and similarly $\|\chi_T^\perp\| = \sqrt{\beta \cdot (1 - \beta) \cdot N}$. \square

Similarly to vertex expansion and edge expansion, a natural question is to what extent the converse holds. That is, if $e(S, T)/ND$ is always “close” to the product of the densities of S and T , then is $\lambda(G)$ necessarily small? This is indeed true:

Theorem 4.16 (Converse to Expander Mixing Lemma). Let G be a D -regular, N -vertex undirected graph. Suppose that for all pairs of disjoint vertex sets S, T , we have $|e(S, T)/(N \cdot D) - \mu(S)\mu(T)| \leq \theta \sqrt{\mu(S)\mu(T)}$ for some $\theta \in [0, 1]$, where $\mu(R) = |R|/N$ for any set R of vertices. Then $\lambda(G) = O(\theta \log(1/\theta))$.

Putting the two theorems together, we see that λ and θ are the same up to a logarithmic factor. Thus, unlike the other connections we have seen, this connection is good for highly expanding graphs (i.e. $\lambda(G)$ close to zero, $\gamma(G)$ close to 1).

4.2 Random Walks on Expanders

From the previous section, we know that one way of characterizing an expander graph G is by having a bound on its second eigenvalue $\lambda(G)$, and in fact there exist constant-degree expanders

where $\lambda(G)$ is bounded by a constant less than 1. From Section 2.4.3, we know that this implies that the random walk on G converges quickly to the uniform distribution. Specifically, a walk of length t started at any vertex ends at ℓ_2 distance at most λ^t from the uniform distribution. Thus after $t = O(\log N)$ steps, the distribution is very close to uniform, e.g. the probability of every vertex is $(1 \pm .01)/N$. Note that, if G has constant degree, the number of random bits invested here is $O(t) = O(\log N)$, which is within a constant factor of optimal; clearly $\log N - O(1)$ random bits are also necessary to sample an almost uniform vertex. Thus, expander walks give a very good tradeoff between the number of random bits invested and the “randomness” of the final vertex in the walk. Remarkably, expander walks give good randomness properties not only for the final vertex in the walk, but also for the *sequence* of vertices traversed in the walk. Indeed, in several ways to be formalized below, this sequence of vertices “behaves” like uniform independent samples of the vertex set.

A canonical application of expander walks is for randomness-efficient error reduction of randomized algorithms: Suppose we have an algorithm with constant error probability, which uses some m random bits. Our goal is to reduce the error to 2^{-k} , with a minimal penalty in random bits and time. Independent repetitions of the algorithm suffers just an $O(k)$ multiplicative penalty in time, but needs $O(km)$ random bits. We have already seen that with pairwise independence we can use just $O(m + k)$ random bits, but the time blows up by $O(2^k)$. Expander graphs let us have the best of both worlds, using just $m + O(k)$ random bits, and increasing the time by only an $O(k)$ factor. Note that for $k = o(m)$, the number of random bits is $(1 + o(1)) \cdot m$, even better than what pairwise independence gives.

The general approach is to consider an expander graph with vertex set $\{0, 1\}^m$, where each vertex is associated with a setting of the random bits. We will choose a uniformly random vertex v_1 and then do a random walk of length $t - 1$, visiting additional vertices v_2, \dots, v_t . (Note that unlike the rapid mixing analysis, here we start at a uniformly random vertex.) This requires m random bits for the initial choice, and $\log D$ for each of the $t - 1$ steps. For every vertex v_i on the random walk, we will run the algorithm with v_i as the setting of the random coins.

First, we consider the special case of randomized algorithms with one-sided error (**RP**). For these, we should accept if at least one execution of the algorithm accepts, and reject otherwise. If the input is a NO instance, the algorithm never accepts, so we also reject. If the input is a YES instance, we want our random walk to hit at least one vertex that makes the algorithm accept. Let B denote the set of “bad” vertices giving coin tosses that make the algorithm reject. By the definition of **RP**, the density of B is at most $1/2$. Thus, our aim is to show that the probability that *all* the vertices in the walk v_1, \dots, v_t are in B vanishes exponentially fast in t , if G is a good expander.

The case $t = 2$ follows from the Expander Mixing Lemma given in the previous section. If we choose a random edge in a graph with spectral expansion $1 - \lambda$, the probability that both endpoints are in a set B is at most $\mu(B)^2 + \lambda \cdot \mu(B)$. So if $\lambda \ll \mu(B)$, then the probability is roughly $\mu(B)^2$, just like two independent random samples. The case of larger t is given by the following theorem.

Theorem 4.17 (Hitting Property of Expander Walks). If G is a regular digraph with spectral expansion $1 - \lambda$, then for any $B \subset V(G)$ of density μ , the probability that a random walk

(V_1, \dots, V_t) of $t - 1$ steps in G starting in a uniformly random vertex V_1 always remains in B is

$$\Pr \left[\bigwedge_{i=1}^t V_i \in B \right] \leq (\mu + \lambda \cdot (1 - \mu))^t$$

Equivalently, a random walk “hits” the complement of B with high probability. Note that if μ and λ are constants less than 1, then the probability of staying in B is $2^{-\Omega(t)}$, completing the analysis of the efficient error-reduction algorithm for **RP**.

Before proving the theorem, we discuss general approaches to analyzing spectral expanders and random walks on them. Typically, the first step is to express the quantities of interest linear-algebraically, involving applications of the random-walk (or adjacency) matrix M to some vectors v . For example, when proving the Expander Mixing Lemma (Lemma 4.15), we expressed the fraction of edges between sets S and T as $\chi_S M \chi_T^t$ (up to some normalization factor). Then we can proceed in one of the two following ways:

Vector Decomposition Decompose the input vector(s) v as $v = v^\parallel + v^\perp$, where $v^\parallel = (\langle v, u \rangle / \langle u, u \rangle)u$ is the component of v in the direction of the uniform distribution u and v^\perp is the component of v orthogonal to u . Then this induces a similar orthogonal decomposition of the output vector vM into

$$vM = (vM)^\parallel + (vM)^\perp = v^\parallel M + v^\perp M,$$

where $v^\parallel M = v^\parallel$ and $\|v^\perp M\| \leq \lambda \cdot \|v^\perp\|$. Thus, from information about how v 's lengths are divided into the uniform and non-uniform components, we deduce information about how vM is divided into the uniform and non-uniform components. This is the approach we took in the proof of the Expander Mixing Lemma.

Matrix Decomposition This corresponds to a different decomposition of the output vector vM that can be expressed in a way that is independent of the decomposition of the input vector v . Specifically, if G has spectral expansion $\gamma = 1 - \lambda$, then

$$vM = v^\parallel + v^\perp M = \gamma \cdot v^\parallel + (\lambda \cdot v^\parallel + v^\perp M) = \gamma \cdot vJ + \lambda \cdot vE = v(\gamma J + \lambda E),$$

where J is the matrix in which every entry is $1/N$ and the error matrix E satisfies $\|vE\| \leq \|v\|$. The advantage of this decomposition is that we can apply it even when we have no information about how v decomposes (only its length). The fact that M is a convex combination of J and E means that we can often treat each of these components separately and then just apply the triangle inequality. However, it is less refined than the vector decomposition approach, and sometimes gives weaker bounds. Indeed, if we used it to prove the Expander Mixing Lemma (without decomposing χ_S and χ_T), we would get a slightly worse error term of $\lambda\sqrt{\mu(S)\mu(T)} + \lambda\mu(S)\mu(T)$.

The Matrix Decomposition Approach can be formalized using the following notion.

Definition 4.18. The (*spectral*) *norm* of an $N \times N$ real matrix M is defined to be

$$\|M\| = \max_{x \in \mathbb{R}^N} \frac{\|xM\|}{\|x\|}$$

(If M is symmetric, then $\|M\|$ equals the *largest* absolute value of any eigenvalue of M .)

Some basic properties of the matrix norm are that $\|cA\| = |c| \cdot \|A\|$, $\|A + B\| \leq \|A\| + \|B\|$, and $\|A \cdot B\| \leq \|A\| \cdot \|B\|$ for every two matrices A, B , and $c \in \mathbb{R}$. Following the discussion above, we have the following lemma:

Lemma 4.19. Let G be a regular digraph on N vertices with random-walk matrix M . Then G has spectral expansion $\gamma = 1 - \lambda$ iff $M = \gamma J + \lambda E$, where J is the $N \times N$ matrix where every entry is $1/N$ (i.e. the random-walk matrix for the complete graph with self-loops) and $\|E\| \leq 1$.

Proof. Suppose that G has spectral expansion γ . Then define $E = (M - \gamma J)/\lambda$. To see that E has norm at most 1, first observe that $uE = (uM - \gamma uJ)/\lambda = (1 - \gamma)u/\lambda = u$. Thus it suffices to show that for every vector v orthogonal to u , the vector vE is orthogonal to u and is of length at most $\|v\|$. Orthogonality follows because vM is orthogonal to u (by regularity of G) and $vJ = 0$. The length bounds follows from $vE = (vM)/\lambda$ and $\|vM\| \leq \lambda\|v\|$ by the spectral expansion of G .

Conversely, suppose that $M = \gamma J + \lambda E$ for some $\|E\| \leq 1$. Then for every vector v orthogonal to u , we have $\|vM\| = \|0 + \lambda vE\| \leq \lambda\|v\|$, and thus G has spectral expansion γ . \square

Intuitively, this lemma says that we can think of a random step on a graph with spectral expansion γ as being a random step on the complete graph with probability γ and “not doing damage” with probability $1 - \gamma$. This intuition would be completely accurate if E were a stochastic matrix, but it is typically not (e.g. it may have negative entries). Still, note that the bound given in Theorem 4.17 exactly matches this intuition: in every step, the probability of remaining in B is at most $\gamma \cdot \mu + \lambda = \mu + \lambda \cdot (1 - \mu)$.

Now we can return to the proof of the theorem.

Proof Theorem 4.17. We need a way to express getting stuck in B linear-algebraically. For that, we define P to be the diagonal matrix with $P_{i,i} = 1$ if $i \in B$ and $P_{i,i} = 0$ otherwise. Thus, the probability a distribution π picks a node in B is $|\pi P|_1$, where $|\cdot|_1$ is the ℓ_1 norm, $|x|_1 = \sum |x_i|$ (which in our case is equal to the sum of the components of the vector, since all values are nonnegative).

Let M be the random-walk matrix of G . The probability distribution for the first vertex V_1 is given by the vector u . Now we can state the following crucial fact:

Claim 4.20. The probability that the random walk stays entirely within B is precisely $|uP(MP)^{t-1}|_1$.

Proof of claim: By induction on ℓ , we show that $(uP(MP)^\ell)_i$ is the probability that the first $\ell + 1$ vertices of the random walk are in B and the $(\ell + 1)$ 'st vertex is i . The base case is $\ell = 0$. If $i \in B$, $(uP)_i = 1/N$; if $i \notin B$, $(uP)_i = 0$. Now assume the hypothesis holds up to some ℓ . Then $(uP(MP)^\ell M)_i$ is the probability that the first $\ell + 1$ vertices of the random walk are in B and the $(\ell + 2)$ 'nd vertex is i (which may or may not be in B). Multiplying by P , we zero out all components for nodes not in B and leave the others unchanged. Thus, we obtain the probability that the first $\ell + 2$ vertices are in B and the $(\ell + 2)$ 'nd vertex is i . \square

To get a bound in terms of the spectral expansion, we will now switch to the ℓ_2 norm. The intuition is that multiplying by M shrinks the component that is perpendicular to u (by expansion) and multiplying by P shrinks the component parallel to u (because it zeroes out some entries). Thus, we should be able to show that the norm $\|MP\|$ is strictly less than 1. Actually, to get the best bound, we note that $uP(MP)^{t-1} = uP(PMP)^{t-1}$, because $P^2 = P$, so we instead bound $\|PMP\|$. Specifically:

Claim 4.21. $\|PMP\| \leq \mu + \lambda \cdot (1 - \mu)$.

Proof of claim: Using the Matrix Decomposition Lemma (Lemma 4.19), we have:

$$\begin{aligned} \|PMP\| &= \|P(\gamma J + \lambda E)P\| \\ &\leq \gamma \cdot \|PJP\| + \lambda \cdot \|PEP\| \\ &\leq \gamma \cdot \|PJP\| + \lambda \end{aligned}$$

Thus, we only need to analyze the case of J , the random walk on the complete graph. Given any vector x , let $y = xP$, so

$$xPJP = yJP = \left(\sum_i y_i \right) uP.$$

Since $\|y\| \leq \|x\|$ and y has at most μN nonzero coordinates, we have

$$\|xPJP\| \leq \left| \sum_i y_i \right| \cdot \|uP\| \leq \left(\sqrt{\mu N} \cdot \|y\| \right) \cdot \sqrt{\frac{\mu}{N}} \leq \mu \cdot \|x\|.$$

Thus,

$$\|PMP\| \leq \gamma \cdot \mu + \lambda = \mu + \lambda \cdot (1 - \mu).$$

□

Using Claims 4.20 and 4.21, the probability of never leaving B in a $(t - 1)$ -step random walk is

$$\begin{aligned} |uP(MP)^{t-1}|_1 &\leq \sqrt{\mu N} \cdot \|uP(MP)^{t-1}\| \\ &\leq \sqrt{\mu N} \cdot \|uP\| \cdot \|PMP\|^{t-1} \\ &\leq \sqrt{\mu N} \cdot \sqrt{\frac{\mu}{N}} \cdot (\mu + \lambda \cdot (1 - \mu))^{t-1} \\ &\leq (\mu + \lambda \cdot (1 - \mu))^t. \end{aligned}$$

□

The hitting properties described above suffice for reducing the error of **RP** algorithms. What about **BPP** algorithms, which have two-sided error? They are handled by the following.

Theorem 4.22 (Chernoff Bound for Expander Walks). Let G be a regular digraph with N vertices and spectral expansion $1 - \lambda$, and let $f : [N] \rightarrow [0, 1]$ be any function. Consider a random walk V_1, \dots, V_t in G from a uniform start vertex V_1 . Then for any $\varepsilon > 0$

$$\Pr \left[\left| \frac{1}{t} \sum_i f(V_i) - \mu(f) \right| \geq \lambda + \varepsilon \right] \leq 2e^{-\Omega(\varepsilon^2 t)}.$$

Note that this is just like the standard Chernoff Bound (Theorem 2.21), except that our additive approximation error increases by $\lambda = 1 - \gamma$. Thus, unlike the Hitting Property we proved above, this bound is only useful when λ is sufficiently small (as opposed to bounded away from 1). This can be achieved by taking the a power of the initial expander, where edges correspond to walks of length t in the original expander; this raises the random-walk matrix and λ to the t 'th power. However, there is a better Chernoff Bound for Expander Walks, where λ does not appear in the approximation error, but the exponent in the probability of error is $\Omega(\gamma\varepsilon^2 t)$ instead of $\Omega(\varepsilon^2 t)$. The bound above suffices in the common case that a small constant approximation error can be tolerated, as in error reduction for **BPP**.

Proof. Let X_i be the random variable $f(V_i)$, and $X = \sum_i X_i$. Just like in the standard proof of the Chernoff Bound (Problem 2.7), we show that the expectation of the moment generating function $e^{rX} = \prod_i e^{rX_i}$ is not much larger than $e^{r\mathbb{E}[X]}$ and apply Markov's Inequality, for a suitable choice of r . However, here the factors e^{rX_i} are not independent, so the expectation does not commute with the product. Instead, we express $\mathbb{E}[e^{rX}]$ linear-algebraically as follows. Define a diagonal matrix P whose (i, i) 'th entry is $e^{rf(i)}$. Then, similarly Claim 4.20 in the proof of the hitting proof above, we observe that

$$\mathbb{E}[e^{rX}] = |uP(MP)^{t-1}|_1 = |u(MP)^t|_1 \leq \sqrt{N} \cdot \|u\| \cdot \|MP\|^t = \|MP\|^t.$$

To see this, we simply note that each cross-term in the matrix product $uP(MP)^{t-1}$ corresponds to exactly one expander walk v_1, \dots, v_t , with a coefficient equal to the probability of this walk times $\prod_i e^{rf(v_i)}$. By the Matrix Decomposition Lemma (Lemma 4.19), we can bound

$$\|MP\| \leq (1 - \lambda) \cdot \|JP\| + \lambda \cdot \|EP\|.$$

Since J simply projects onto the uniform direction, we have

$$\begin{aligned} \|JP\|^2 &= \frac{\|uP\|^2}{\|u\|^2} \\ &= \frac{\sum_v (e^{rf(v)}/N)^2}{\sum_v (1/N)^2} \\ &= \frac{1}{N} \cdot \sum_v e^{2rf(v)} \\ &= \frac{1}{N} \cdot \sum_v (1 + 2rf(v) + O(r^2)) \\ &= 1 + 2r\mu + O(r^2) \end{aligned}$$

for $r \leq 1$, and thus

$$\|JP\| = \sqrt{1 + 2r\mu + O(r^2)} = 1 + r\mu + O(r^2).$$

For the error term, we have

$$\|EP\| \leq \|P\| \leq e^r = 1 + r + O(r^2).$$

Thus,

$$\|MP\| \leq (1 - \lambda) \cdot (1 + r\mu + O(r^2)) + \lambda \cdot (1 + r + O(r^2)) \leq 1 + (\mu + \lambda)r + O(r^2),$$

and we have

$$\mathbb{E}[e^{rX}] \leq (1 + (\mu + \lambda)r + O(r^2))^t \leq e^{(\mu + \lambda)rt + O(r^2t)}.$$

By Markov's Inequality,

$$\Pr[X \geq (\mu + \lambda + \varepsilon) \cdot t] \leq e^{-\varepsilon rt + O(r^2t)} = e^{-\Omega(\varepsilon^2 t)},$$

if we set $r = \varepsilon/c$ for a large enough constant c . By applying the same analysis to the function $1 - f$, we see that $\Pr[X \leq (\mu - \lambda - \varepsilon)t] = e^{-\Omega(\varepsilon^2 t)}$, and this establishes the theorem. \square

We now summarize the properties that expander walks give us for randomness-efficient error reduction and sampling.

For reducing the error of a **BPP** algorithm from $1/3$ to 2^{-k} , we can apply Theorem 4.22 with $\lambda = \varepsilon = 1/12$, so that a walk of length $t = O(k)$ suffices. If the original **BPP** algorithm used m random bits and the expander is of constant degree (which is possible with $\lambda = 1/12$), then the number of random bits needed is only $m + O(k)$. Comparing with previous methods for error reduction, we have:

	Number of Repetitions	Number of Random Bits
Independent Repetitions	$O(k)$	$O(km)$
Pairwise Independent Repetitions	$O(2^k)$	$O(k + m)$
Expander Walks	$O(k)$	$m + O(k)$

For SAMPLING, where we are given an oracle to a function $f : \{0, 1\}^m \rightarrow [0, 1]$ and we want to approximate $\mu(f)$ to within an additive error of ε , we can apply Theorem 4.22 with error $\varepsilon/2$ and $\lambda = \varepsilon/2$. The needed expander can be obtained by taking an $O(\log(1/\varepsilon))$ 'th power of a constant-degree expander, yielding the following bounds:

	Number of Samples	Number of Random Bits
Truly Random Sample	$O((1/\varepsilon^2) \cdot \log(1/\delta))$	$O(m \cdot (1/\varepsilon^2) \cdot \log(1/\delta))$
Pairwise Independent Samples	$O((1/\varepsilon^2) \cdot (1/\delta))$	$O(m + \log(1/\varepsilon) + \log(1/\delta))$
Expander Walks	$O((1/\varepsilon^2) \cdot \log(1/\delta))$	$m + O(\log(1/\delta) \cdot (\log(1/\varepsilon)/\varepsilon^2))$

The $\log(1/\varepsilon)$ factor in the number of random bits comes because we took an $O(\log(1/\varepsilon))$ 'th power of a constant-degree expander and thus spend $O(\log(1/\varepsilon))$ random bits for each step on the expander. This is actually not necessary and comes from the slightly weaker Chernoff Bound we proved. In any case, note that expander walks have a much better dependence on the error probability δ in the number of samples (as compared to pairwise independence), but have a worse dependence on the approximation error ε in the number of random bits. Problem 4.5 shows how to combine these two samplers to achieve the best of both.

Similarly to pairwise independence, the sampling algorithm based on expander walks is actually an averaging sampler in the sense of Definition 3.29:

Theorem 4.23 (Expander-Walk Sampler). For every $m \in \mathbb{N}$ and $\delta, \varepsilon \in [0, 1]$, there is a (δ, ε) averaging sampler $\text{Samp} : \{0, 1\}^n \rightarrow (\{0, 1\}^m)^t$ using $n = m + O(\log(1/\delta) \cdot \log(1/\varepsilon)/\varepsilon^2)$ random bits and $t = O((1/\varepsilon^2) \cdot \log(1/\delta))$ samples.

The sampling algorithm of Problem 4.5 that combines expander walks and pairwise independence, however, is *not* an averaging sampler, and it is an open problem to achieve similar parameters with an explicit averaging sampler:

Open Problem 4.24. Give an explicit construction of a (δ, ε) averaging sampler $\text{Samp} : \{0, 1\}^n \rightarrow (\{0, 1\}^m)^t$ that uses $n = O(m + \log(\delta) + \log(1/\varepsilon))$ random bits and $t = O((1/\varepsilon^2) \cdot \log(1/\delta))$ samples.

Before we end this section, we make an important remark: we have not actually given an efficient algorithm for randomness-efficient error reduction (or an explicit averaging sampler)! Our algorithm assumes an expander graph of exponential size, namely 2^m where m is the number of random bits used by the algorithm. Generating such a graph at random would use far too many coins. Even generating it deterministically would not suffice, since we would have to write down an exponential-size object. In the following section, we will see how to define an explicit expander graph without writing it down in its entirety, and efficiently do random walks in such a graph.

4.3 Explicit Constructions

As discussed in previous sections, expander graphs have numerous applications in theoretical computer science. (See also the Chapter Notes and Exercises.) For some of these applications, it may be acceptable to simply choose the graph at random, as we know that a random graph will be a good expander with high probability. For many applications, however, this simple approach does not suffice. Some reasons are the following (in increasing order of significance):

- We may not want to tolerate the error probability introduced by the (unlikely) event that the graph is not an expander. To deal with this, we could try checking that the graph is an expander. Computing the expansion of a given graph is **NP**-hard for most of the combinatorial measures (e.g. vertex expansion or edge expansion), but the spectral expansion can be computed to high precision in time polynomial in the size of the graph (as it is just an eigenvalue computation). As we saw, spectral expansion does yield estimates on vertex expansion and edge expansion (but cannot give optimal expansion in these measures).
- Some of the applications of expanders (like the one from the previous section) are for reducing the amount of randomness needed for certain tasks. Thus choosing the graph at random defeats the purpose.
- A number of the applications require *exponentially large* expander graphs, and thus we cannot even write down a randomly chosen expander. For example, for randomness-efficient error reduction of randomized algorithms, we need an expander on 2^m nodes where m is the number of random bits used by the algorithm.

From a more philosophical perspective, finding explicit constructions is a way of developing and measuring our understanding of these important combinatorial objects.

A couple of alternatives for defining explicit constructions of expanders on N nodes are:

Mildly Explicit: Construct a complete representation of the graph in time $\text{poly}(N)$.

Fully Explicit: Given a node $u \in [N]$ and $i \in [D]$, where D is the degree of the expander, compute the i^{th} neighbor of u in time $\text{poly}(\log N)$.

Consider the randomness-efficient error reduction application discussed in the previous section, in which we performed a random walk on an expander graph with exponentially many nodes. Mild explicitness is insufficient for this application, as the desired expander graph is of exponential size, and hence cannot be even entirely stored, let alone constructed. But full explicitness is perfectly suited for efficiently conducting a random walk on a huge graph. So now our goal is the following:

Goal: Devise a fully explicit construction of an infinite family $\{G_i\}$ of D -regular graphs with spectral expansion at least γ , where D and $\gamma > 0$ are constants independent of i .

We remark that we would also like the set $\{N_i\}$, where N_i is the number of vertices in G_i , to be not too sparse, so that the family of graphs $\{G_i\}$ has graphs of size close to any desired size.

4.3.1 Algebraic Constructions

Here we mention a few known explicit constructions that are of interest because of their simple description, the parameters achieved, and/or the mathematics that goes into their analysis. We will not prove the expansion properties of any of these constructions (but will rather give a different explicit construction in the subsequent sections).

Construction 4.25 (discrete torus expanders). Let $G = (V, E)$ be the graph with vertex set $V = \mathbb{Z}_M \times \mathbb{Z}_M$, and edges from each node (x, y) to the nodes (x, y) , $(x + 1, y)$, $(x, y + 1)$, $(x, x + y)$, $(-y, x)$, where all arithmetic is modulo M .

This is a fully explicit 5-regular digraph with $N = M^2$ nodes and spectral expansion $\gamma = \Omega(1)$. It can be made undirected by adding a reverse copy of each edge. We refer to these as “discrete torus” expanders because \mathbb{Z}_M^2 can be viewed as a discrete version of the real torus, namely $[0, 1]^2$ with arithmetic modulo 1. The expansion of these graphs was originally proved using group representation theory, but later proofs for similar discrete-torus expanders were found that only rely on Fourier analysis over the torus.

Construction 4.26 (p -cycle with inverse chords). This is the graph $G = (V, E)$ with vertex set $V = \mathbb{Z}_p$ and edges that connect each node x with the nodes: $x + 1$, $x - 1$ and x^{-1} (where all arithmetic is mod p and we define 0^{-1} to be 0).

This graph is only mildly explicit since we do not know how to construct n -bit primes deterministically in time $\text{poly}(n)$ (though Cramér’s conjecture in Number Theory would imply that we can do so by simply checking the first $\text{poly}(n)$ n -bit numbers). The proof of expansion relies on the “Selberg 3/16 Theorem” from number theory.

Construction 4.27 (Ramanujan graphs). $G = (V, E)$ is a graph with vertex set $V = \mathbb{F}_q \cup \{\infty\}$, the finite field of prime order q s.t. $q \equiv 1 \pmod{4}$ plus one extra node representing infinity. The edges in this graph connect each node z with all z' of the form

$$z' = \frac{(a_0 + ia_1)z + (a_2 + ia_3)}{(-a_2 + ia_3)z + (a_0 - ia_1)}$$

for $a_0, a_1, a_2, a_3 \in \mathbb{Z}$ such that $a_0^2 + a_1^2 + a_2^2 + a_3^2 = p$, a_0 is odd and positive, and a_1, a_2, a_3 are even, for some fixed prime $p \neq q$ such that $p \equiv 1 \pmod{4}$, q is a square modulo p , and $i \in \mathbb{F}_q$ such that $i^2 = -1 \pmod{q}$.

The degree of the graph is the number of solutions to the equation $a_0^2 + a_1^2 + a_2^2 + a_3^2 = p$, which turns out to be $D = p + 1$, and it has $\lambda(G) \leq 2\sqrt{D-1}/D$, so it is an *optimal* spectral expander. (See Theorems 4.11 and 4.12, and note that this bound is even better than we know for random graphs, which have an additive $o(1)$ term in the spectral expansion.) These graphs are also only mildly explicit, again due to the need to find the prime q .

These are called Ramanujan Graphs because the proof of their spectral expansion relies on results in number theory concerning the “Ramanujan Conjectures.” Subsequently, the term *Ramanujan graphs* came to refer to any infinite family of graphs with optimal spectral expansion $\gamma \geq 1 - 2\sqrt{D-1}/D$.

4.3.2 Graph Operations

The explicit construction of expanders given in the next section will be an iterative one, where we start with a “constant size” expander H and repeatedly apply graph operations to get bigger expanders. The operations that we apply should increase the number of nodes in the graph, while keeping the degree and the second eigenvalue λ bounded. We’ll see three operations, each improving one property while paying a price on the others; however, combined together, they yield the desired expander. It turns out that this approach for constructing expanders will also be useful in derandomizing the logspace algorithm for UNDIRECTED S-T CONNECTIVITY, as we will see in Section 4.4.

The following concise notation will be useful to keep track of each of the parameters:

Definition 4.28. An (N, D, γ) -graph is a D -regular digraph on N vertices with spectral expansion γ .

4.3.2.1 Squaring

Definition 4.29 (Squaring of Graphs). If $G = (V, E)$ is a D -regular digraph, then $G^2 = (V, E')$ is a D^2 -regular digraph on the same vertex set, where the (i, j) ’th neighbor of a vertex x is the j ’th neighbor of the i ’th neighbor of x . In particular, a random step on G^2 consists of two random steps on G .

Lemma 4.30. If G is a $(N, D, 1 - \lambda)$ -graph, then G^2 is a $(N, D^2, 1 - \lambda^2)$ -graph.

Namely, the degree deteriorates by squaring, while the spectral expansion is improved from $\gamma = 1 - \lambda$ to $\gamma' = 1 - \lambda^2 = 2\gamma - \gamma^2$.

Proof. The effect of squaring on the number of nodes N and the degree D is immediate from the definition. For the spectral expansion, note that if M is the random-walk matrix for G , then M^2 is the random-walk matrix for G^2 . So for any vector $x \perp u$,

$$\|xM^2\| \leq \lambda \cdot \|xM\| \leq \lambda^2 \cdot \|x\|.$$

□

4.3.2.2 Tensoring

The next operation we consider increases the size of the graph at the price of increasing the degree.

Definition 4.31 (Tensor Product of Graphs). Let $G_1 = (V_1, E_1)$ be D_1 -regular and $G_2 = (V_2, E_2)$ be D_2 -regular. Then their *tensor product* is the D_1D_2 -regular graph $G_1 \otimes G_2 = (V_1 \times V_2, E)$, where the (i_1, i_2) 'th neighbor of a vertex (x_1, x_2) is (y_1, y_2) , where y_b is the i_b 'th neighbor of x_b in G_b . That is, a random step on $G_1 \otimes G_2$ consists of a random step on G_1 in the first component and a random step on G_2 in the second component.

Often this operation is simply called the “product” of G_1 and G_2 , but we use “tensor product” to avoid confusion with squaring and to reflect its connection with the standard tensor products in linear algebra:

Definition 4.32 (Tensor Products of Vectors and Matrices). Let $x \in \mathbb{R}^{N_1}, y \in \mathbb{R}^{N_2}$, then their *tensor product* is the vector $z = x \otimes y \in \mathbb{R}^{N_1N_2}$ where $z_{ij} = x_i y_j$.

Similarly, for matrices $A = (a_{ij}) \in \mathbb{R}_{N_1 \times N_1}, B = (b_{ij}) \in \mathbb{R}_{N_2 \times N_2}$, their *tensor product* is the matrix $C = A \otimes B \in \mathbb{R}_{N_1N_2 \times N_1N_2}$ where $C = (c_{ij,i'j'})$ for $c_{ij,i'j'} = a_{ii'} b_{jj'}$.

A few comments on the tensor operation:

- A random walk on a tensor graph $G_1 \otimes G_2$ is equivalent to taking two independent random walks on G_1 and G_2 .
- For vectors $x \in \mathbb{R}^{N_1}, y \in \mathbb{R}^{N_2}$ that are probability distributions (i.e. nonnegative vectors with ℓ_1 norm 1), their tensor product $x \otimes y$ is a probability distribution on $[N_1] \times [N_2]$ where the two components are independently distributed according to x and y , respectively.
- $(x \otimes y)(A \otimes B) = (xA) \otimes (yB)$ for every $x \in \mathbb{R}^{N_1}, y \in \mathbb{R}^{N_2}$, and in fact $A \otimes B$ is the unique matrix with this property.
- Not all vectors $z \in \mathbb{R}^{N_1N_2}$ are decomposable as $x \otimes y$ for $x \in \mathbb{R}^{N_1}$ and $y \in \mathbb{R}^{N_2}$. Nevertheless, the set of all decomposable tensors $x \otimes y$ spans $\mathbb{R}^{N_1N_2}$.

- If M_1, M_2 are the random-walk matrices for graphs G_1, G_2 respectively, then the random-walk matrix for the graph $G_1 \otimes G_2$ is

$$M_1 \otimes M_2 = (I_{N_1} \otimes M_2)(M_1 \otimes I_{N_2}) = (M_1 \otimes I_{N_2})(I_{N_1} \otimes M_2),$$

where I_N denotes the $N \times N$ identity matrix. That is, we can view a random step on $G_1 \otimes G_2$ as being a random step on the G_1 component followed by one on the G_2 component or vice-versa.

The effect of tensoring on expanders is given by the following:

Lemma 4.33. If G_1 is an (N_1, D_1, γ_1) -graph and G_2 is an (N_2, D_2, γ_2) -graph, then $G_1 \otimes G_2$ is an $(N_1 N_2, D_1 D_2, \min\{\gamma_1, \gamma_2\})$ -graph.

In particular, if $G_1 = G_2$, then the number of nodes improves, the degree deteriorates, and the spectral expansion remains unchanged.

Proof. As usual, we write $\gamma_1 = 1 - \lambda_1$, $\gamma_2 = 1 - \lambda_2$; then our goal is to show that $G_1 \otimes G_2$ has spectral expansion $1 - \max\{\lambda_1, \lambda_2\}$. The intuition is as follows. We can think of the vertices of $G_1 \otimes G_2$ as being partitioned into N_1 “clouds,” each consisting of N_2 vertices, where cloud v_1 contains all vertices of the form (v_1, \cdot) . Thus, any probability distribution (V_1, V_2) on the vertices (v_1, v_2) of $G_1 \otimes G_2$ can be thought of as picking a cloud v_1 according to the marginal distribution² V_1 and then picking the vertex v_2 within the cloud v_1 according to the conditional distribution $V_2|_{V_1=v_1}$. If the overall distribution on pairs is far from uniform, then either

- (1) The marginal distribution V_1 on the clouds must be far from uniform, *or*
- (2) the conditional distributions $V_2|_{V_1=v_1}$ within the clouds must be far from uniform.

When we take a random step, the expansion of G_1 will bring us closer to uniform in Case 1 and the expansion of G_2 will bring us closer to uniform in Case 2.

One way to prove the bound in the case of undirected graphs is to use the fact that the eigenvalues of $M_1 \otimes M_2$ are all the products of eigenvalues of M_1 and M_2 , so the largest magnitude is $1 \cdot 1$ and the next largest is bounded by either $\lambda_1 \cdot 1$ or $1 \cdot \lambda_2$. Instead, we use the Vector Decomposition Method to give a proof that matches the intuition more closely and is a good warm-up for the analysis of the zig-zag product in the next section. Given any vector $x \in \mathbb{R}^{N_1 N_2}$ that is orthogonal to $u_{N_1 N_2}$, we can decompose x as $x = x^{\parallel} + x^{\perp}$ where x^{\parallel} is a multiple of u_{N_2} on each cloud of size N_2 and x^{\perp} is orthogonal to u_{N_2} on each cloud. Note that $x^{\parallel} = y \otimes u_{N_2}$, where $y \in \mathbb{R}^{N_1}$ is orthogonal to u_{N_1} (because $x^{\parallel} = x - x^{\perp}$ is orthogonal to $u_{N_1 N_2}$). If we think of x as the nonuniform component of a probability distribution, then x^{\parallel} and x^{\perp} correspond to the two cases in the intuition above.

For the first case, we have

$$x^{\parallel} M = (y \otimes u_{N_2})(M_1 \otimes M_2) = (y M_1) \otimes u_{N_2}.$$

The expansion of G_1 tells us that M_1 shrinks y by a factor of λ_1 , and thus $\|x^{\parallel} M\| \leq \lambda_1 \cdot \|x^{\parallel}\|$. For the second case, we write

$$x^{\perp} M = x^{\perp} (I_{N_1} \otimes M_2)(M_1 \otimes I_{N_2}).$$

²For two jointly distributed random variables (X, Y) , the *marginal distribution* of X is simply the distribution of X alone, ignoring information about Y .

The expansion of G_2 tells us that M_2 will shrink x^\perp by a factor of λ_2 on each cloud, and thus $I_{N_1} \otimes M_2$ will shrink x^\perp by the same factor. The subsequent application of $M_1 \otimes I_{N_2}$ cannot increase the length (being the random-walk matrix for a regular graph, albeit a disconnected one). Thus, $\|x^\perp M\| \leq \lambda_2 \|x^\perp\|$.

Finally, we argue that $x^\parallel M$ and $x^\perp M$ are orthogonal. Note that $x^\parallel M = (yM_1) \otimes u_{N_2}$ is a multiple of u_{N_2} on every cloud. Thus it suffices to argue that x^\perp remains orthogonal to u_{N_2} on every cloud after we apply M . Applying $(I_{N_1} \otimes M_2)$ retains this property (because applying M_2 preserves orthogonality to u_{N_2} , by regularity of G_2) and applying $(M_1 \otimes I_{N_2})$ retains this property because it assigns each cloud a linear combination of several other clouds (and a linear combination of vectors orthogonal to u_{N_2} is also orthogonal to u_{N_2}).

Thus,

$$\begin{aligned} \|xM\|^2 &= \|x^\parallel M\|^2 + \|x^\perp M\|^2 \\ &\leq \lambda_1^2 \cdot \|x^\parallel\|^2 + \lambda_2^2 \cdot \|x^\perp\|^2 \\ &\leq \max\{\lambda_1, \lambda_2\}^2 \cdot (\|x^\parallel\|^2 + \|x^\perp\|^2) \\ &= \max\{\lambda_1, \lambda_2\}^2 \cdot \|x\|^2, \end{aligned}$$

as desired. □

4.3.2.3 The Zig-Zag Product

Of the two operations we have seen, one (squaring) improves expansion and one (tensoring) increases size, but both have the deleterious effect of increasing the degree. Now we will see a third operation that decreases the degree, without losing too much in the expansion. By repeatedly applying these three operations, we will be able to construct arbitrarily large expanders while keeping both the degree and expansion constant.

Let G be an (N_1, D_1, γ_1) expander and H be a (D_1, D_2, γ_2) expander. The *zig-zag product* of G and H , denoted $G \mathbb{Z} H$, will be defined as follows. The nodes of $G \mathbb{Z} H$ are the pairs (u, i) where $u \in V(G)$ and $i \in V(H)$. The edges of $G \mathbb{Z} H$ will be defined so that a random step on $G \mathbb{Z} H$ corresponds to a step on G , but using a random step on H to choose the edge in G . (This is the reason why we require the number of vertices in H to be equal to the degree of G .) A step in $G \mathbb{Z} H$ will therefore involve a step to a random neighbor in H and then a step in G to a neighbor whose index is equal to the label of the current node in H . Intuitively, a random walk on a “good” expander graph H should generate choices that are sufficiently random to produce a “good” random walk on G . One problem with this definition is that it is not symmetric. That is, the fact that you can go from (u, i) to (v, j) does not mean that you can go from (v, j) to (u, i) . We correct this by adding another step in H after the step in G . In addition to allowing us to construct undirected expander graphs, this extra step will also turn out to be important for the expansion of $G \mathbb{Z} H$.

More formally,

Definition 4.34 (Zig-zag Product). Let G be an D_1 -regular digraph on N_1 vertices, and H a D_2 -regular digraph on D_1 vertices. Then $G \mathbb{Z} H$ is a graph whose vertices are pairs $(u, i) \in [N_1] \times [D_1]$. For $a, b \in [D_2]$, the (a, b) 'th neighbor of a vertex (u, i) is the vertex (v, j) computed as follows:

- (1) Let i' be the a 'th neighbor of i in H .

- (2) Let v be the i' 'th neighbor of u in G , so $e = (u, v)$ is the i' 'th edge leaving u . Let j' be such that e is the j' 'th edge entering v in G . (In an undirected graph, this simply means that u is the j' 'th neighbor of v .)
- (3) Let j be the b' 'th neighbor of j' in H .

Note that the graph $G \otimes H$ depends on how the edges leaving and entering each vertex of G are numbered. Thus it is best thought of as an operation on labelled graphs. (This is made more explicit in Section 4.3.3 via the notion of an “edge-rotation map.”) Nevertheless, the bound we will prove on expansion holds regardless of the labelling:

Theorem 4.35. If G is a (N_1, D_1, γ_1) -graph, and H is a (D_1, D_2, γ_2) -graph then $G \otimes H$ is a $(N_1 D_1, D_2^2, \gamma = \gamma_1 \cdot \gamma_2^2)$ -graph. In particular, if $\gamma_1 = 1 - \lambda_1$ and $\gamma_2 = 1 - \lambda_2$, then $\gamma = 1 - \lambda$ for $\lambda \leq \lambda_1 + 2\lambda_2$.

G should be thought of as a big graph and H as a small graph, where D_1 is a large constant and D_2 is a small constant. Note that the number of nodes D_1 in H is required to equal the degree of G . Observe that when $D_1 > D_2^2$ the degree is reduced by the zig-zag product.

There are two different intuitions underlying the expansion of the zig-zag product:

- Given an initial distribution (U, I) on the vertices of $G_1 \otimes G_2$ that is far from uniform, there are two extreme cases, just as in the intuition for the tensor product.³ Either
 - (1) All the (conditional) distributions $I|_{U=u}$ within the clouds are far from uniform,
 - or*
 - (2) All the (conditional) distributions $I|_{U=u}$ within the clouds of size D_1 are uniform (in which case the marginal distribution U on the clouds must be far from uniform).

In Case 1, the first H -step $(U, I) \mapsto (U, I')$ already brings us closer to the uniform distribution, and the other two steps cannot hurt (as they are steps on regular graphs). In Case 2, the first H -step has no effect, but the G -step $(U, I') \mapsto (V, J')$ has the effect of making the marginal distribution on clouds closer to uniform, i.e. V is closer to uniform than U . But note that the joint distribution (V, J') isn't actually any closer to the uniform distribution on the vertices of $G_1 \otimes G_2$ because the G -step is a permutation. Still, if the marginal distribution V on clouds is closer to uniform, then the conditional distributions within the clouds $J'|_{V=v}$ must have become further from uniform, and thus the second H -step $(V, J') \mapsto (V, J)$ brings us closer to uniform. This leads to a proof by *Vector Decomposition*, where we decompose any vector x that is orthogonal to uniform into components x^{\parallel} and x^{\perp} , where x^{\parallel} is uniform on each cloud, and x^{\perp} is orthogonal to uniform on each cloud. This approach gives the best known bounds on the spectral expansion of the zig-zag product, but it can be a bit messy since the two components generally do not remain orthogonal after the steps of the zig-zag product (unlike the case of the tensor product, where we were able to show that $x^{\parallel}M$ is orthogonal to $x^{\perp}M$).

³ Here we follow our convention of using capital letters to denote random variables corresponding to the lower-case values in Definition 4.34.

- The second intuition is to think of the expander H as behaving “similarly” to the complete graph on D_1 vertices (with self-loops). In the case that H equals the complete graph, then it is easy to see that $G \mathbb{Z} H = G \otimes H$. Thus it is natural to apply *Matrix Decomposition*, writing the random-walk matrix for an arbitrary expander H as a convex combination of the random-walk matrix for the complete graph and an error matrix. This gives a very clean analysis, but slightly worse bounds than the Vector Decomposition Method.

We now proceed with the formal proof, following the Matrix Decomposition approach.

Proof of Theorem 4.35. Let A , B , and M be the random-walk matrices for G_1 , G_2 , and $G_1 \mathbb{Z} G_2$, respectively. We decompose M into the product of three matrices, corresponding to the three steps in the definition of $G_1 \mathbb{Z} G_2$'s edges. Let \tilde{B} be the transition matrix for taking a random G_2 -step on the second component of $[N_1] \times [D_1]$, i.e. $\tilde{B} = I_{N_1} \otimes B$, where I_{N_1} is the $N_1 \times N_1$ identity matrix. Let \hat{A} be the permutation matrix corresponding to the G_1 -step. That is, $\hat{A}_{(u,i),(v,j)}$ is 1 iff (u, v) is the i 'th edge leaving u and the j 'th edge entering v . By the definition of $G_1 \mathbb{Z} G_2$, we have $M = \tilde{B} \hat{A} \tilde{B}$.

By the Matrix Decomposition Lemma (Lemma 4.19), $B = \gamma_2 J + (1 - \gamma_2) E$, where every entry of J equals $1/D_1$ and E has norm at most 1. Then $\tilde{B} = \gamma_2 \tilde{J} + (1 - \gamma_2) \tilde{E}$, where $\tilde{J} = I_{N_1} \otimes J$ and $\tilde{E} = I_{N_1} \otimes E$ has norm at most 1.

This gives

$$M = \left(\gamma_2 \tilde{J} + (1 - \gamma_2) \tilde{E} \right) \hat{A} \left(\gamma_2 \tilde{J} + (1 - \gamma_2) \tilde{E} \right) = \gamma_2^2 \tilde{J} \hat{A} \tilde{J} + (1 - \gamma_2^2) F,$$

where we take $(1 - \gamma_2^2) F$ to be the sum of the three terms involving \tilde{E} ; noting that their norms sum to at most $(1 - \gamma_2^2)$, we see that F has norm at most 1. Now, the key observation is that $\tilde{J} \hat{A} \tilde{J} = A \otimes J$.

Thus,

$$M = \gamma_2^2 \cdot A \otimes J + (1 - \gamma_2^2) F,$$

and thus

$$\begin{aligned} \lambda(M) &\leq \gamma_2^2 \cdot \lambda(A \otimes J) + (1 - \gamma_2^2) \\ &\leq \gamma_2^2 \cdot (1 - \gamma_1) + (1 - \gamma_2^2) \\ &= 1 - \gamma_1 \gamma_2^2, \end{aligned}$$

as desired. □

4.3.3 The Expander Construction

As a first attempt for constructing a family of expanders, we construct an infinite family G_1, G_2, \dots of graphs utilizing only the squaring and the zig-zag operations:

Construction 4.36 (Mildly Explicit Expanders). Let H be a $(D^4, D, 7/8)$ -graph (e.g., as constructed in Problem 4.8), and define:

$$\begin{aligned} G_1 &= H^2 \\ G_{t+1} &= G_t^2 \mathbb{Z} H \end{aligned}$$

Proposition 4.37. For all t , G_t is a $(D^{4t}, D^2, 1/2)$ -graph.

Proof. By induction on t .

Base Case: by the definition of H and Lemma 4.30, $G_1 = H^2$ is a $(D^4, D^2, 1 - \lambda_0^2)$ -graph and $\lambda_0^2 \leq 1/2$.

Induction Step: First note that $G_t^2 \otimes H$ is well-defined because $\deg(G_t^2) = \deg(G_t)^2 = (D^2)^2 = \#\text{nodes}(H)$. Then,

$$\begin{aligned} \deg(G_{t+1}) &= \deg(H)^2 = D^2 \\ \#\text{nodes}(G_{t+1}) &= \#\text{nodes}(G_t^2) \cdot \#\text{nodes}(H) = N_t \cdot D^4 = D^{4t} D^4 = D^{4(t+1)} \\ \lambda(G_{t+1}) &\leq \lambda(G_t)^2 + 2\lambda(H) \leq (1/2)^2 + 2 \cdot (1/8) = 1/2 \end{aligned}$$

□

Now, we recursively bound the time to compute neighbors in G_t . Actually, due to the way the G -step in the zig-zag product is defined, we bound the time to compute the *edge-rotation map* $(u, i) \mapsto (v, j)$, where the i 'th edge leaving u equals the j 'th edge entering v . Denote by $\text{time}(G_t)$ the time required for one evaluation of the edge-rotation map for G_t . This requires two evaluations of the edge-rotation map for G_{t-1} (the squaring requires two applications, while the zig-zag part does not increase the number of applications), plus time $\text{poly}(\log N_t)$ for manipulating strings of length $O(\log N_t)$. Therefore,

$$\begin{aligned} \text{time}(G_t) &= 2 \cdot \text{time}(G_{t-1}) + \text{poly}(\log N_t) \\ &= 2^t \cdot \text{poly}(\log N_t) \\ &= N_t^{\Theta(1)}, \end{aligned}$$

where the last equality holds because $N_t = D^{4t}$ for a constant D . Thus, this construction is only mildly explicit.

We remedy the above difficulty by using tensoring to make the sizes of the graphs grow more quickly:

Construction 4.38 (Fully Explicit Expanders). Let H be a $(D^8, D, 7/8)$ -graph, and define:

$$\begin{aligned} G_1 &= H^2 \\ G_{t+1} &= (G_t \otimes G_t)^2 \otimes H \end{aligned}$$

In this family of graphs, the number of nodes grows doubly exponentially $N_t \approx c^{2^t}$, while the computation time grows only exponentially as before. Namely,

$$\text{time}(G_t) = 4^t \cdot \text{poly}(\log N_t) = \text{poly}(\log N_t).$$

We remark that the above family is rather sparse, i.e. the numbers in $\{N_t\}$ are far apart. To overcome this shortcoming, we can amend the above definition to have

$$G_t = (G_{\lceil t/2 \rceil} \otimes G_{\lfloor t/2 \rfloor})^2 \otimes H.$$

Now $N_t = D^{8t}$, so given a number N , we can find a graph G_t in the family whose size is at most $D^8 \cdot N = O(N)$. Moreover, the construction remains fully explicit because $\text{time}(G_t) = O(\text{time}(G_{\lceil t/2 \rceil}) + \text{time}(G_{\lfloor t/2 \rfloor})) = \text{poly}(t)$. Thus we have established:

Theorem 4.39. There is a constant $D \in \mathbb{N}$ such that for every $t \in \mathbb{N}$, there is a fully explicit expander graph G_t with degree D , spectral expansion $1/2$, and $N_t = D^{4t}$ nodes.

Consequently, the randomness-efficient error-reduction and averaging sampler based on expander walks can be made explicit:

Corollary 4.40. If a language L has a **BPP** algorithm with error probability at most $1/3$ that uses $m(n)$ random bits on inputs of length n , then for every polynomial $k(n)$, L has a **BPP** algorithm with error probability at most $2^{-k(n)}$ that uses $m(n) + O(k(n))$ random bits.

Corollary 4.41. There is an explicit averaging sampler achieving the parameters of Theorem 4.23.

4.3.4 Open Problems

As we have seen, spectral expanders such as those in Theorem 4.39 are also vertex expanders (Theorem 4.6 and Corollary 4.10) and edge expanders (Theorem 4.14), but these equivalences do not extend to optimizing the various expansion measures.

As mentioned in Section 4.3.1, there are known explicit constructions of optimal spectral expanders, namely Ramanujan graphs. However, unlike the expanders of Theorem 4.39, those constructions rely on deep results in number theory. The lack of a more elementary construction seems to signify a limitation in our understanding of expander graphs.

Open Problem 4.42. Give an explicit “combinatorial” construction of constant-degree expander graphs G with $\lambda(G) \leq 2\sqrt{D-1}/D$ (or even $\lambda(G) = O(1/\sqrt{D})$), where D is the degree.

For vertex expansion, it is known how to construct *bipartite* (or directed) expanders with constant left-degree (or out-degree) D and expansion $(1-\varepsilon) \cdot D$ for an arbitrarily small constant ε (see Chapter 6), but achieving the optimal expansion of $D - O(1)$ (cf., Theorem 4.4) or constructing undirected vertex expanders with high expansion remains open.

Open Problem 4.43. For an arbitrarily large constant D , give an explicit construction of bipartite $(\Omega(N), D - c)$ vertex expanders with N vertices on each side and left-degree D , where c is a universal constant independent of D .

Open Problem 4.44. For an arbitrarily small constant $\varepsilon > 0$, give an explicit construction of *undirected* $(\Omega(N), (1-\varepsilon)D)$ vertex expanders with N vertices and constant degree D that depends only on ε .

We remark that while Open Problem 4.43 refers to balanced bipartite graphs (i.e. ones with the same number of vertices on each side), the imbalanced case is also interesting and important. (See Problems 4.10, 5.5 and Open Problems 5.36, 6.35.)

4.4 UNDIRECTED S-T CONNECTIVITY in Deterministic Logspace

Recall the UNDIRECTED S-T CONNECTIVITY problem: given an undirected graph G and two vertices s, t , decide whether there is a path from s to t . In Section 2.4, we saw that this problem can be solved in randomized logspace (**RL**). Here we will see how we can use expanders and the operations above to solve this problem in deterministic logspace (**L**).

The algorithm is based on the following two ideas:

- UNDIRECTED S-T CONNECTIVITY can be solved in logspace on constant-degree expander graphs. More precisely, it is easy on constant-degree graphs where every connected component is promised to be an expander (i.e. has spectral expansion bounded away from 1): we can try all paths of length $O(\log N)$ from s in logarithmic space; this works because expanders have logarithmic diameter. (See Problem 4.2.)
- The same operations we used to construct an infinite expander family above can also be used to turn *any* graph into an expander (in logarithmic space). Above, we started with a constant-sized expander and used various operations to build larger and larger expanders. There, the goal was to increase the size of the graph (which was accomplished by tensoring and/or zig-zag), while preserving the degree and the expansion (which was accomplished by zig-zag and squaring, which made up for losses in these parameters). Here, we want to improve the expansion (which will be accomplished by squaring), while preserving the degree (as will be handled by zig-zag) and ensuring the graph remains of polynomial size (so we will not use tensoring).

Specifically, the algorithm is as follows.

Algorithm 4.45 (UNDIRECTED S-T CONNECTIVITY in L).

Input: An undirected graph G with N edges and vertices s and t .

- (1) Let H be a fixed $(D^4, D, 3/4)$ graph for some constant D .
 - (2) Reduce (G, s, t) to (G_0, s_0, t_0) , where G_0 is a D^2 -regular graph in which every connected component is nonbipartite and s_0 and t_0 are connected in G_0 iff s and t are connected in G .
 - (3) For $k = 1, \dots, \ell = O(\log N)$, define:
 - (a) Let $G_k = G_{k-1}^2 \otimes H$
 - (b) Let s_k and t_k be any two vertices in the “clouds” of G_k corresponding to s_{k-1} and t_{k-1} , respectively. (Note that if s_k and t_k are connected in G_k , then s_{k-1} and t_{k-1} are connected in G_{k-1} .)
 - (4) Try all paths of length $O(\log N)$ in G_ℓ from s_ℓ and accept if any of them visit t_ℓ .
-

We will discuss how to implement this algorithm in logspace later, and first analyze its correctness. Let C_k be the connected component of G_k containing s_k . Observe that C_k is a connected component of $C_{k-1}^2 \otimes H$; below we will show that $C_{k-1}^2 \otimes H$ is connected and hence $C_k = C_{k-1}^2 \otimes H$. Since C_0 is undirected, connected, and nonbipartite, we have $\gamma(C_0) \geq 1/\text{poly}(N)$ by Theorem 2.53. We will argue that in each iteration the spectral gap increases by a constant factor, and thus after $O(\log N)$ iterations we have an expander.

By Lemma 4.30, we have

$$\gamma(C_{k-1}^2) \geq 2 \cdot \gamma(C_{k-1}) \cdot (1 - \gamma(C_{k-1})/2) \approx 2\gamma(C_k)$$

for small $\gamma(C_{k-1})$. By Theorem 4.35, we have

$$\begin{aligned} \gamma(C_{k-1}^2 \otimes H) &\geq \gamma(H)^2 \cdot \gamma(C_{k-1}^2) \\ &\geq \left(\frac{3}{4}\right)^2 \cdot 2 \cdot \gamma(C_{k-1}) \cdot (1 - \gamma(C_{k-1})/2) \\ &\geq \min \left\{ \frac{35}{32} \cdot \gamma(C_{k-1}), \frac{1}{18} \right\} \end{aligned}$$

where the last inequality is obtained by considering whether $\gamma(C_{k-1}) \leq 1/18$ or $\gamma(C_{k-1}) > 1/18$. In particular, $C_{k-1}^2 \otimes H$ is connected, so we have $C_k = C_{k-1}^2 \otimes H$ and

$$\gamma(C_k) \geq \min \left\{ \frac{35}{32} \cdot \gamma(C_{k-1}), \frac{1}{18} \right\}.$$

Thus, after $\ell = O(\log N)$ iterations, we must have $\gamma(C_\ell) \geq 1/18$. Moreover, observe that the number of vertices N_ℓ in G_ℓ is at most $N_0 \cdot (D^4)^\ell = \text{poly}(N)$, so considering paths of length $O(\log N)$ will suffice to decide s - t connectivity in G_ℓ .

To show that the algorithm can be implemented in logarithmic space, we argue that the edge-rotation map of each G_k can be computed with only $O(1)$ more space than the edge-rotation map of G_{k-1} , so that G_ℓ requires space $O(\log N) + O(\ell) = O(\log N)$. Since the inductive claim here refers to sublogarithmic differences of space (indeed $O(1)$ space) and sublogarithmic space is model-dependent (even keeping a pointer into the input requires logarithmic space), we will refer to a specific model of computation in establishing it. (The final result, that UNDIRECTED S-T CONNECTIVITY is in \mathbf{L} , is, however, model-independent because logspace computations in any reasonable computational model can be simulated by logspace computations in any other reasonable model.) Formally, let $\text{space}(G_k)$ denote the workspace needed to compute the edge-rotation map of G_k on a multi-tape Turing machine with the following input/output conventions:

- Input Description:
 - Tape 1 (read-only): Contains the initial input graph G , with the head at the leftmost position of the tape.
 - Tape 2 (read-write): Contains the input pair (v, i) , where v is a vertex of G_i and $i \in [D^2]$ is an index of the a neighbor on a *read-write* tape, with the head at the *rightmost* position of i . The rest of the tape may contain additional data.
 - Tapes 3+ (read-write): Blank worktapes with the head at the leftmost position.
- Output Description:
 - Tape 1: The head should be returned to the leftmost position.
 - Tape 2: In place of (v, i) , it should contain the output (w, j) where w is the i 'th neighbor of v and v is the j 'th neighbor of w . The head should be at the rightmost position of j and the rest of the tape should remain unchanged from its state at the beginning of the computation.

- Tapes 3+ (read-write): Are returned to the blank state with the heads at the leftmost position.

With these conventions, it is not difficult to argue that $\text{space}(G_0) = O(\log N)$, and $\text{space}(G_k) = \text{space}(G_{k-1}) + O(1)$. For the latter, we first argue that $\text{space}(G_{k-1}^2) = \text{space}(G_{k-1}) + O(1)$, and then that $\text{space}(G_{k-1}^2 \otimes H) = \text{space}(G_{k-1}^2) + O(1)$. For G_{k-1}^2 , we are given a triple $(v, (i_1, i_2))$ on tape 2, with the head on the rightmost position of i_2 , and both i_1 and i_2 are elements of $[D^2]$ (and thus of constant size). We move the head left to the rightmost position of i_1 , compute the edge-rotation map of G_{k-1} on (v, i_1) so that tape 2 now contains (w, j_1, i_2) . Then we swap j_1 and i_2 , and run the edge-rotation map of G_{k-1} on (w, i_2) to get (w, j_2, j_1) , and move the head to the rightmost position of j_1 , completing the rotation. For $G_{k-1}^2 \otimes H$, we are given a tuple $((v, i), (a_1, a_2))$, where v is a vertex of G_{k-1}^2 , i is a vertex of H (equivalently, an edge-label for G_{k-1}^2), and a_1, a_2 are edge labels for H . Evaluating the rotation map requires two evaluations of the rotation map for H (both of which are “constant-size” operations) and one evaluation of the rotation map of G_{k-1}^2 .

Thus we have proven:

Theorem 4.46. **UNDIRECTED S-T CONNECTIVITY** is in **L**.

We remark that proving $\mathbf{RL} = \mathbf{L}$ in general remains open. The best deterministic simulation known for \mathbf{RL} is essentially $\mathbf{L}^{3/2} = \mathbf{DSPACE}(\log^{3/2} n)$, which makes beautiful use of known pseudorandom generators for logspace computation. (Unfortunately, we do not have space to cover this line of work in this survey.) Historically, improved derandomizations for **UNDIRECTED S-T CONNECTIVITY** have inspired improved derandomizations of \mathbf{RL} (and vice-versa). Since Theorem 4.46 is still quite recent (2005), there is a good chance that we have not yet exhausted the ideas in it.

Open Problem 4.47. Show that $\mathbf{RL} \subset \mathbf{L}^c$ for some constant $c < 3/2$.

Another open problem is the construction of *universal traversal sequences* — fixed walks of polynomial length that are guaranteed to visit all vertices in any connected undirected regular graph of a given size. (See Example 3.8 and Open Problem 3.9.) Using the ideas from the algorithm above, it is possible to obtain logspace-constructible, polynomial-length universal traversal sequences for all regular graphs that are *consistently labelled* in the sense that no pair of distinct vertices have the same i 'th neighbor for any $i \in [D]$. For general labellings, the best known universal traversal sequences are of length $N^{O(\log N)}$ (and are constructible in space $O(\log^2 N)$).

Open Problem 4.48 (Open Problem 3.9, restated). Give an explicit construction of universal traversal sequences of polynomial length for arbitrarily labelled undirected graphs (or even for an arbitrary labelling of the complete graph).

We remark that handling general labellings (for “pseudorandom walk generators” rather than universal traversal sequences) seems to be the main obstacle in extending the techniques of Theorem 4.46 to prove $\mathbf{RL} = \mathbf{L}$. (See the Chapter Notes and References.)

4.5 Exercises

Problem 4.1 (Bipartite vs. Nonbipartite Expanders). Show that constructing bipartite expanders is equivalent to constructing (standard, nonbipartite) expanders. That is, show how given an explicit construction of one of the following, you can obtain an explicit construction of the other:

- (1) D -regular $(\alpha N, A)$ expanders on N vertices for infinitely many N , where $\alpha > 0$, $A > 1$, and D are constants independent of N .
- (2) D -regular (on both sides) $(\alpha N, A)$ bipartite expanders with N vertices on each side for infinitely many N , where $\alpha > 0$, $A > 1$, and D are constants independent of N .

(Your transformations need not preserve the constants.)

Problem 4.2 (More Combinatorial Consequences of Spectral Expansion). Let G be a graph on N vertices with spectral expansion $\gamma = 1 - \lambda$. Prove that:

- (1) The *independence number* $\alpha(G)$ is at most $(\lambda/(1 + \lambda))N$, where $\alpha(G)$ is defined to be the size of the largest independent set, i.e. subset S of vertices s.t. there are no edges with both endpoints in S .
- (2) The *chromatic number* $\chi(G)$ is at least $(1 + \lambda)/\lambda$, where $\chi(G)$ is defined to be the smallest number of colors for which the vertices of G can be colored s.t. all pairs of adjacent vertices have different colors.
- (3) The *diameter* of G is $O(\log_{1/\lambda} N)$.

Recall that computing $\alpha(G)$ and $\chi(G)$ exactly are **NP**-complete problems. However, the above shows that for expanders, nontrivial bounds on these quantities can be computed in polynomial time.

Problem 4.3 (Limits on Vertex Expansion). This problem and the next one give limits on the vertex and spectral expansion that can be achieved as a function of the degree D . Both bounds are proved by relating the expansion of an arbitrary D -regular graph G by that of the infinite D -regular tree T_D (where every vertex has one parent and $D - 1$ children), which is in some sense the “best possible” D -regular expander.

- (1) Show that if a D -regular digraph G is a (K, A) expander, then T_D is a (K, A) expander.
 - (2) Show that for every $D \in \mathbb{N}$, there are infinitely many $K \in \mathbb{N}$ such that T_D is not a $(K, D - 1 + 2/K)$ expander.
 - (3) Deduce that for constant $D \in \mathbb{N}$ and $\alpha > 0$, if a D -regular, N -vertex digraph G is an $(\alpha N, A)$ vertex expander, then $A \leq D - 1 + o(1)$, where the $o(1)$ term vanishes as $N \rightarrow \infty$ (and D, α are held constant).
-

Problem 4.4 (Limits on Spectral Expansion). Let G be a D -regular undirected graph and T_D be the infinite D -regular tree (as in Problem 4.3). For a graph H and $\ell \in \mathbb{N}$, let $p_\ell(H)$ denote the probability that if we choose a random vertex v in H and do a random walk of length 2ℓ , we end back at vertex v .

- (1) Show that $p_\ell(G) \geq p_\ell(T_D) \geq C_\ell \cdot (D-1)^\ell / D^{2\ell}$, where C_ℓ is the ℓ 'th Catalan number, which equals the number of properly parenthesized strings in $\{(,)\}^{2\ell}$ — strings where no prefix has more $)$'s than $($'s.
- (2) Show that $N \cdot p_\ell(G) \leq 1 + (N-1) \cdot \lambda(G)^{2\ell}$. (Hint: use the fact that the trace of a matrix equals the sum of its eigenvalues.)
- (3) Using the fact that $C_\ell = \binom{2\ell}{\ell} / (\ell+1)$, prove that

$$\lambda(G) \geq \frac{2\sqrt{D-1}}{D} - o(1),$$

where the $o(1)$ term vanishes as $N \rightarrow \infty$ (and D is held constant).

Problem 4.5 (Near-Optimal Sampling). (1) Describe an algorithm for SAMPLING that tosses $O(m + \log(1/\varepsilon) + \log(1/\delta))$ coins, makes $O((1/\varepsilon^2) \cdot \log(1/\delta))$ queries to a function $f : \{0, 1\}^m \rightarrow [0, 1]$, and estimates $\mu(f)$ to within $\pm\varepsilon$ with probability at least $1 - \delta$. (Hint: use expander walks to generate several sequences of coin tosses for the pairwise-independent averaging sampler, and compute the answer via a “median of averages”.)

- (2) Give an explicit (δ, ε) hitting sampler (see Problem 3.9) $\text{Samp} : \{0, 1\}^n \rightarrow (\{0, 1\}^m)^t$ that tosses $n = O(m + \log(1/\varepsilon) + \log(1/\delta))$ coins and generates $t = O((1/\varepsilon) \cdot \log(1/\delta))$ samples.

It turns out that these bounds on the randomness and query/sample complexities are each optimal up to constant factors (for most parameter settings of interest).

Problem 4.6 (Error Reduction For Free*). Show that if a problem has a **BPP** algorithm with constant error probability, then it has a **BPP** algorithm with error probability $1/n$ that uses *exactly* the same number of random bits.

Problem 4.7 (Vertex Expanders vs. Hitting Samplers). Here we will see that hitting samplers (defined in Problem 3.9) are equivalent to a variant of vertex expanders, where we only require that for (left-)sets S of size *exactly* K , there are at least $A \cdot K$ neighbors. We call such graphs $(= K, A)$ *vertex expanders* and will revisit them in the next chapter (Definition 5.32).

Given a bipartite multigraph with neighbor function $\Gamma : [N] \times [D] \rightarrow [M]$, we can obtain a sampler $\text{Samp} : [N] \rightarrow [M]^D$ by setting $\text{Samp}(x)_y = \Gamma(x, y)$. Conversely, every such sampler gives rise to a bipartite multigraph. Prove that Samp is a (δ, ε) hitting sampler if and only if Γ is an $(= K, A)$ vertex expander for $K = \lfloor \delta N \rfloor + 1$ and $A = (1 - \varepsilon)M/K$.

Thus, bipartite vertex expanders and hitting samplers are *equivalent*. However, the typical settings of parameters for the two objects are very different. For example, in vertex expanders, a primary goal is usually to maximize the expansion factor A , but $A \cdot K$ may be significantly smaller than M . In samplers, $AK = (1 - \varepsilon)M$ is usually taken to be very close to M , but $A = (1 - \varepsilon)M/\delta N$ may even be smaller than 1. Similarly, the most common setting of expanders takes $K/N \approx \delta$ to be a constant, whereas in samplers it is often thought of as vanishingly small.

Problem 4.8 (A “Constant-Sized” Expander). (1) Let \mathbb{F} be a finite field. Consider a graph G with vertex set \mathbb{F}^2 and edge set $\{(a, b), (c, d) : ac = b + d\}$. That is, we connect vertex (a, b) to all points on the line $y = ax - b$. Prove that G is $|\mathbb{F}|$ -regular and $\lambda(G) \leq 1/\sqrt{|\mathbb{F}|}$. (Hint: consider G^2 .)

(2) Show that if $|\mathbb{F}|$ is sufficiently large (but still constant), then by applying appropriate operations to G , we can obtain a base graph for the expander construction given in Section 4.3.3, i.e. a $(D^8, D, 7/8)$ graph for some constant D .

Problem 4.9 (The Replacement Product). Given a D_1 -regular graph G_1 on N_1 vertices and a D_2 -regular graph G_2 on D_1 vertices, consider the following graph $G_1 \textcircled{F} G_2$ on vertex set $[N_1] \times [D_1]$: vertex (u, i) is connected to (v, j) iff (a) $u = v$ and (i, j) is an edge in G_2 , or (b) v is the i 'th neighbor of u in G_1 and u is the j 'th neighbor of v . That is, we “replace” each vertex v in G_1 with a copy of G_2 , associating each edge incident to v with one vertex of G_2 .

- (1) Prove that there is a function g such that if G_1 has spectral expansion $\gamma_1 > 0$ and G_2 has spectral expansion $\gamma_2 > 0$ (and both graphs are undirected), then $G_1 \textcircled{F} G_2$ has spectral expansion $g(\gamma_1, \gamma_2, D_2) > 0$. (Hint: Note that $(G_1 \textcircled{F} G_2)^3$ has $G_1 \textcircled{Z} G_2$ as a subgraph.)
 - (2) Show how to convert an explicit construction of constant-degree (spectral) expanders into an explicit construction of degree 3 (spectral) expanders.
 - (3) Without using Theorem 4.14, prove an analogue of Part 1 for edge expansion. That is, there is a function h such that if G_1 is an $(N_1/2, \varepsilon_1)$ edge expander and G_2 is a $(D_1/2, \varepsilon_2)$ edge expander, then $G_1 \textcircled{F} G_2$ is a $(N_1 D_1/2, h(\varepsilon_1, \varepsilon_2, D_2))$ edge expander, where $h(\varepsilon_1, \varepsilon_2, D_2) > 0$ if $\varepsilon_1, \varepsilon_2 > 0$. (Hint: given any set S of vertices of $G_1 \textcircled{F} G_2$, partition S into the clouds that are more than “half-full” and those that are not.)
 - (4) Prove that the functions $g(\gamma_1, \gamma_2, D_2)$ and $h(\varepsilon_1, \varepsilon_2, D_2)$ must depend on D_2 , by showing that $G_1 \textcircled{F} G_2$ cannot be a $(N_1 D_1/2, \varepsilon)$ edge expander if $\varepsilon > 1/(D_2 + 1)$ and $N_1 \geq 2$.
-

Problem 4.10 (Unbalanced Vertex Expanders and Data Structures). Consider a $(K, (1 - \varepsilon)D)$ bipartite vertex expander G with N left vertices, M right vertices, and left degree D .

- (1) For a set S of left vertices, a $y \in N(S)$ is called a *unique neighbor* of S if y is incident to exactly one edge from S . Prove that every left-set S of size at most K has at least $(1 - 2\varepsilon)D|S|$ unique neighbors.

- (2) For a set S of size at most $K/2$, prove that at most $|S|/2$ vertices outside S have at least δD neighbors in $N(S)$, for $\delta = O(\varepsilon)$.

Now we'll see a beautiful application of such expanders to data structures. Suppose we want to store a small subset S of a large universe $[N]$ such that we can test membership in S by probing just 1 bit of our data structure. A trivial way to achieve this is to store the characteristic vector of S , but this requires N bits of storage. The hashing-based data structures mentioned in Section 3.5.3 only require storing $O(|S|)$ words, each of $O(\log N)$ bits, but testing membership requires reading an entire word (rather than just one bit.)

Our data structure will consist of M bits, which we think of as a $\{0, 1\}$ -assignment to the right vertices of our expander. This assignment will have the following property.

Property II: For all left vertices x , all but a $\delta = O(\varepsilon)$ fraction of the neighbors of x are assigned the value $\chi_S(x)$ (where $\chi_S(x) = 1$ iff $x \in S$).

- (3) Show that if we store an assignment satisfying Property II, then we can probabilistically test membership in S with error probability δ by reading just one bit of the data structure.
- (4) Show that an assignment satisfying Property II exists provided $|S| \leq K/2$. (Hint: first assign 1 to all of S 's neighbors and 0 to all its nonneighbors, then try to correct the errors.)

It turns out that the needed expanders exist with $M = O(K \log N)$ (for any constant ε), so the size of this data structure matches the hashing-based scheme while admitting (randomized) 1-bit probes. However, note that such bipartite vertex expanders do *not* follow from explicit spectral expanders as given in Theorem 4.39, because the latter do not provide vertex expansion beyond $D/2$ nor do they yield highly imbalanced expanders (with $M \ll N$) as needed here. But in Chapter 5, we will see how to explicitly construct expanders that are quite good for this application (specifically, with $M = K^{1.01} \cdot \text{polylog} N$).

4.6 Chapter Notes and References

A detailed coverage of expander graphs and their applications in theoretical computer science is given by Hoory, Linial, and Wigderson [HLW]. Applications in pure mathematics are surveyed by Lubotzky [Lub2].

The first papers on expander graphs appeared in conferences on telephone networks. Specifically, Pinsker [Pin] proved that random graphs are good expanders, and used these to demonstrate the existence of graphs called “concentrators.” Bassalygo [Bas] improved Pinsker’s results, in particular giving the general tradeoff between the degree D , expansion factor A , and set density α mentioned after Theorem 4.4. The first computer science application of expanders (and “superconcentrators”) came in an approach by Valiant [Val1] to proving circuit lower bounds. An early and striking algorithmic application was the $O(\log n)$ -depth sorting network by Ajtai, Komlós, and Szemerédi [AKS1], which also illustrated the usefulness of expanders for derandomization. An exciting recent application of expanders is Dinur’s new proof of the PCP Theorem [Din].

The fact that spectral expansion implies vertex expansion and edge expansion was shown by Tanner [Tan] (for vertex expansion) and Alon and Milman [AM] (for edge expansion). The converses

are discrete analogues of Cheeger’s Inequality for Riemannian manifolds [Che1], and various forms of these were proven by Alon [Alo1] (for vertex expansion), Jerrum and Sinclair [JS] (for edge expansion in undirected graphs and, more generally, conductance in reversible Markov chains), and Mihail [Mih] (for edge expansion in regular digraphs and conductance in non-reversible Markov chains).

The “Ramanujan” upper bound on spectral expansion given by Theorem 4.11 was proven by Alon and Boppana (see [Alo1, Nil]). Theorem 4.12, stating that random graphs are asymptotically Ramanujan, was conjectured by Alon [Alo1], but was only proven recently by Friedman [Fri]. Kahale [Kah] proved that Ramanujan graphs have vertex expansion roughly $D/2$ for small sets.

Forms of the Expander Mixing Lemma date back to Alon and Chung [AC2], who considered the number of edges between a set and its complement (i.e. $T = V \setminus S$). The converse to the Expander Mixing Lemma (Theorem 4.16) is due to Bilu and Linial [BL]. For more on quasirandomness, see [CGW, AS1] for the case of dense graphs and [CG1, CG2] for sparse graphs.

The sampling properties of random walks on expanders were analyzed in a series of works starting with Ajtai, Komlós, and Szemerédi [AKS2]. The hitting bound of Theorem 4.17 is due to Kahale [Kah], and the Chernoff Bound for expander walks (cf., Theorem 4.22) is due to Gillman [Gil3]. Our proof of the Chernoff Bound is inspired by that of Healy [Hea], who also provides some other variants and generalizations. The **RP** version of Problem 4.6 is due to Karp, Pippenger, and Sipser [KPS], who initiated the study of randomness-efficient error reduction of randomized algorithms. It was generalized to **BPP** in [CW1]. The equivalence of hitting samplers and bipartite vertex expanders from Problem 4.7 is due to Sipser [Sip2]. Problem 4.5 is due to Bellare, Goldreich, and Goldwasser [BGG]; matching lower bounds for sampling were given by Canetti, Even, and Goldreich [CEG]. Open Problem 4.24 was posed by Bellare and Rompel [BR].

Construction 4.25 is due to Margulis [Mar1], and was the first explicit construction of constant-degree expanders. Gabber and Galil [GG] (see also [JM]) gave a much more elementary proof of expansion for similar expanders, which also provided a specific bound on the spectral expansion (unlike Margulis’ proof). Construction 4.26 is variant of a construction of Lubotzky, Phillips, and Sarnak. (See [Lub1, Thm. 4.42], from which the expansion of Construction 4.26 can be deduced.) Ramanujan graphs (Construction 4.27) were constructed independently by Lubotzky, Phillips, and Sarnak [LPS] and Margulis [Mar2]. For more on Ramanujan graphs and the mathematical machinery that goes into their analysis, see the books [Sar, Lub1, DSV].

The zig-zag product and the expander constructions of Section 4.3.3 are due to Reingold, Vadhan, and Wigderson [RVW2]. Our analysis of the zig-zag product is from [RTV2], which in turn builds on [RV], who used matrix decomposition (Lemma 4.19) for analyzing other graph products. Earlier uses of graph products in constructing expanders include the use of the tensor product in [Tan]. Problem 4.9, on the replacement product, is from [RVW2, RTV2], and can be used in place of the zig-zag product in both the expander constructions and the **UNDIRECTED S-T CONNEC-TIVITY** algorithm (Algorithm 4.45). Independently of [RVW2], Martin and Randall [MR] proved a “decomposition theorem” for Markov chains that implies a better bound on the spectral expansion of the replacement product.

There has been substantial progress on giving a combinatorial construction of Ramanujan graphs (Open Problem 4.42). Bilu and Linial [BL] give a mildly explicit construction achieving $\lambda(G) = \tilde{O}(\sqrt{D})$, Ben-Aroya and Ta-Shma [BT] give a fully explicit construction achieving $\lambda(G) = D^{1/2+o(1)}$,

and Spielman et al. [BSS] give a mildly explicit construction of a *weighted* graph achieving $\lambda(G) = O(\sqrt{D})$.

Constant-degree bipartite expanders with expansion $(1 - \varepsilon) \cdot D$ have been constructed by Capalbo et al. [CRVW], based on a variant of the zig-zag product for “randomness condensers.” (See Section 6.3.5.) Alon and Capalbo [AC1] have made progress on Open Problem 4.44 by giving an explicit construction of undirected constant-degree “unique-neighbor” expanders (see Problem 4.10).

The deterministic logspace algorithm for UNDIRECTED S-T CONNECTIVITY (Algorithm 4.45) is due to Reingold [Rei2]. The result that $\mathbf{RL} \subset \mathbf{L}^{3/2}$ is due to Saks and Zhou [SZ1], with an important ingredient being Nisan’s pseudorandom generator for space-bounded computation [Nis2]. Based on Algorithm 4.45, explicit polynomial-length universal traversal sequences for “consistently labelled” regular digraphs, as well as “pseudorandom walk generators” for such graphs, were constructed in [Rei2, RTV2]. (See also [RV].) In [RTV2], it is shown that pseudorandom walk generators for arbitrarily labelled regular digraphs would imply $\mathbf{RL} = \mathbf{L}$. The best known explicit construction of a full-fledged universal traversal sequence is due to Nisan [Nis2], has length $n^{O(\log n)}$, and can be constructed in time $n^{O(\log n)}$ and space $O(\log^2 n)$. (See Section 8.2.1 for more on the derandomization of \mathbf{RL} .)

Problem 4.8, Part 1 is a variant of a construction of Alon [Alo2]; Part 2 is from [RVW2]. The results of Problem 4.2 are from [Hof, Lov2, AM, Chu]. The result of Problem 4.10, on bit-probe data structures for set membership, is due to Buhrman, Miltersen, Radhakrishnan, and Venkatesan [BMRV].

5

List-Decodable Codes

The field of *coding theory* is motivated by the problem of communicating reliably over noisy channels — where the data sent over the channel may come out corrupted on the other end, but we nevertheless want the receiver to be able to correct the errors and recover the original message. There is a vast literature studying aspects of this problem from the perspectives of electrical engineering (communications and information theory), computer science (algorithms and complexity), and mathematics (combinatorics and algebra). In this survey, we are interested in codes as “pseudorandom objects.” In particular, a generalization of the notion of an error-correcting code yields a framework that we will use to unify all of the main pseudorandom objects covered in this survey (averaging samplers, expander graphs, randomness extractors, list-decodable codes, pseudorandom generators).

5.1 Definitions and Existence

5.1.1 Definition.

The approach to communicating over a noisy channel is to restrict the data we send to be from a certain set of strings that can be easily disambiguated (even after being corrupted).

Definition 5.1. A q -ary code is a multiset¹ $\mathcal{C} \subset \Sigma^{\hat{n}}$, where Σ is an *alphabet* of size q . Elements of \mathcal{C} are called *codewords*. We define the following key parameters:

- \hat{n} is the *block length*.
- $n = \log |\mathcal{C}|$ is the *message length*.
- $\rho = n/(\hat{n} \cdot \log |\Sigma|)$ is the (*relative*) *rate* of the code.

An *encoding function* for \mathcal{C} is a bijective mapping $\text{Enc}: \{1, \dots, |\mathcal{C}|\} \rightarrow \mathcal{C}$. Given such an encod-

¹Traditionally, codes are defined to be sets rather than multisets. However, the generality afforded by multisets will allow the connections we see later to be stated more cleanly. In the case \mathcal{C} is a multiset, the condition that a mapping $\text{Enc}: \{1, \dots, N\} \rightarrow \mathcal{C}$ is bijective means that for every $c \in \mathcal{C}$, $|\text{Enc}^{-1}(c)|$ equals the multiplicity of c in \mathcal{C} .

ing function, we view the elements of $\{1, \dots, |\mathcal{C}|\}$ as *messages*. When $n = \log |\mathcal{C}|$ is an integer, we often think of messages as strings in $\{0, 1\}^n$.

We view \mathcal{C} and Enc as being essentially the same object (with Enc merely providing a “labelling” of codewords), with the former being useful for studying the combinatorics of codes and the latter for algorithmic purposes.

We remark that our notation differs from the standard notation in coding theory in several ways. Typically in coding theory, the input alphabet is taken to be the same as the output alphabet (rather than $\{0, 1\}$ and Σ , respectively), the blocklength is denoted n , and the message length (over Σ) is denoted k and is referred to as the rate. Our notational choices are made in part to facilitate stating connections to other pseudorandom objects covered in this survey.

So far, we haven’t talked at all about the error-correcting properties of codes. Here we need to specify two things: the model of errors (as introduced by the noisy channel) and the notion of a successful recovery.

For the errors, the main distinction is between *random errors* and *worst-case errors*. For random errors, one needs to specify a stochastic model of the channel. The most basic one is the *binary symmetric channel* (over alphabet $\Sigma = \{0, 1\}$), where each bit is flipped independently with probability δ . People also study more complex channel models, but as usual with stochastic models, there is always the question of how well the theoretical model correctly captures the real-life distribution of errors. We, however, will focus on *worst-case errors*, where we simply assume that fewer than a δ fraction of symbols have been changed. That is, when we send a codeword $c \in \Sigma^{\hat{n}}$ over the channel, the received word $r \in \Sigma^{\hat{n}}$ differs from c in fewer than $\delta \hat{n}$ places. Equivalently, c and r are close in Hamming distance:

Definition 5.2 (Hamming distance). For two strings $x, y \in \Sigma^{\hat{n}}$, their (*relative*) *Hamming distance* $d_H(x, y)$ equals $\Pr_i[x_i \neq y_i]$. The *agreement* is defined to be $\text{agr}(x, y) = 1 - d_H(x, y)$.

For a string $x \in \Sigma^{\hat{n}}$ and $\delta \in [0, 1]$, the (*open*) *Hamming ball* of radius δ around x is the set $B(x, \delta)$ of strings $y \in \Sigma^{\hat{n}}$ such that $d_H(x, y) < \delta$.

For the notion of a successful recovery, the traditional model requires that we can uniquely decode the message from the received word (in the case of random errors, this need only hold with high probability). Our main focus will be on a more relaxed notion which allows us to produce a small list of candidate messages. As we will see, the advantage of such list decoding is that it allows us to correct a larger fraction of errors.

Definition 5.3. Let $\text{Enc}: \{0, 1\}^n \rightarrow \Sigma^{\hat{n}}$ be an encoding algorithm for a code \mathcal{C} . A δ -*decoding* algorithm for Enc is a function $\text{Dec}: \Sigma^{\hat{n}} \rightarrow \{0, 1\}^n$ such that for every $m \in \{0, 1\}^n$ and $r \in \Sigma^{\hat{n}}$ satisfying $d_H(\text{Enc}(m), r) < \delta$, we have $\text{Dec}(r) = m$. If such a function Dec exists, we call the code δ -*decodable*.

A (δ, L) -*list-decoding* algorithm for Enc is a function $\text{Dec}: \Sigma^{\hat{n}} \rightarrow (\{0, 1\}^n)^L$ such that for every $m \in \{0, 1\}^n$ and $r \in \Sigma^{\hat{n}}$ satisfying $d_H(\text{Enc}(m), r) < \delta$, we have $m \in \text{Dec}(r)$. If such a function Dec exists, we call the code (δ, L) -*list-decodable*.

Note that a δ -decoding algorithm is the same as a $(\delta, 1)$ -list-decoding algorithm. It is not hard to see that, if we do not care about computational efficiency, the existence of such decoding algorithms depends only on the combinatorics of the set of codewords.

Definition 5.4. The (relative) minimum distance of a code $\mathcal{C} \subset \Sigma^{\hat{n}}$ equals $\min_{x \neq y \in \mathcal{C}} d_H(x, y)$.²

Proposition 5.5. Let $\mathcal{C} \subset \Sigma^{\hat{n}}$ be a code with any encoding function Enc.

- (1) For $\delta \hat{n} \in \mathbb{N}$, Enc is δ -decodable iff its minimum distance is at least $2\delta - 1/\hat{n}$.
- (2) Enc is (δ, L) -list-decodable iff for every $r \in \Sigma^{\hat{n}}$, we have $|B(r, \delta) \cap \mathcal{C}| \leq L$.

Proof. Item 2 follows readily from the definitions.

For Item 1, first note that Enc is δ -decodable iff there is no received word $r \in \Sigma^{\hat{n}}$ at distance less than δ from two codewords c_1, c_2 . (Such an r should decode to both c_1 and c_2 , which is impossible for a unique decoder.) If the code has minimum distance at least $2\delta - 1/\hat{n}$, then c_1 and c_2 disagree in at least $2\delta\hat{n} - 1$ positions, which implies that r disagrees with one of them in at least $\delta\hat{n}$ positions. (Recall that $\delta\hat{n}$ is an integer by hypothesis.) Conversely, if the code has minimum distance smaller than $2\delta - 1/\hat{n}$, then there are two codewords that disagree in at most $2\delta\hat{n} - 2$ positions, and we can construct r that disagrees with each in at most $\delta\hat{n} - 1$ positions. \square

Because of the factor of 2 in Item 1, unique decoding is only possible at distances up to $1/2$, whereas we will see that list decoding is possible at distances approaching 1 (with small lists).

The main goals in constructing codes are to have infinite families of codes (e.g. for every message length n) in which we:

- Maximize the fraction δ of errors correctible (e.g. constant independent of n and \hat{n}).
- Maximize the rate ρ (e.g. a constant independent of n and \hat{n}).
- Minimize the alphabet size q (e.g. a constant, ideally $q = 2$).
- Keep the list size L relatively small (e.g. a constant or $\text{poly}(n)$).
- Have computationally efficient encoding and decoding algorithms.

In particular, coding theorists are very interested in obtaining the optimal tradeoff between the constants δ and ρ with efficiently encodable and decodable codes.

5.1.2 Existential Bounds

Like expanders, the existence of very good codes can be shown using the probabilistic method. The bounds will be stated in terms of the q -ary entropy functions, so we begin by defining those.

Definition 5.6. For $q, \hat{n} \in \mathbb{N}$ and $\delta \in [0, 1]$, we define $H_q(\delta, \hat{n}) \in [0, 1]$ to be such that $|B(x, \delta)| = q^{H_q(\delta, \hat{n}) \cdot \hat{n}}$ for $x \in \Sigma^{\hat{n}}$, where Σ is an alphabet of size q .

We also define the q -ary entropy function $H_q(\delta) = \delta \cdot \log_q((q-1)/\delta) + (1-\delta) \cdot \log_q(1/(1-\delta))$.

The reason we use similar notation for $H_q(\delta, \hat{n})$ and $H_q(\delta)$ is Part 1 of the following:

²If any codeword appears in \mathcal{C} with multiplicity greater than 1, then the minimum distance is defined to be zero.

Proposition 5.7. For every $q \in \mathbb{N}$, $\delta \in (0, 1 - 1/q)$, $\varepsilon \in [0, 1/2]$,

- (1) $\lim_{\hat{n} \rightarrow \infty} H_q(\delta, \hat{n}) = H_q(\delta)$.
 - (2) $H_q(\delta) \leq H_2(\delta)/\log q + \delta$.
 - (3) $H_2(1/2 - \varepsilon) = 1 - \Theta(\varepsilon^2)$.
-

Now we state the bounds for random error-correcting codes.

Theorem 5.8. (1) For all $\hat{n}, q \in \mathbb{N}$ and $\delta \in (0, 1 - 1/q)$, there exists a q -ary code of block length \hat{n} , minimum distance at least δ , and rate at least $\rho = 1 - H_q(\delta, \hat{n})$.

(2) For all integers $\hat{n}, q, L \in \mathbb{N}$ and $\delta \in (0, 1 - 1/q)$, there exists a (δ, L) -list-decodable q -ary code of block length \hat{n} and rate at least $\rho = 1 - H_q(\delta, \hat{n}) - 1/(L + 1)$.

Proof. (1) Pick the codewords c_1, \dots, c_N in sequence ensuring that c_i is at distance at least δ from c_1, \dots, c_{i-1} . The union of the Hamming balls of radius δ around c_1, \dots, c_{i-1} contains at most $(i - 1) \cdot q^{H_q(\delta, \hat{n}) \cdot \hat{n}} < (N - 1) \cdot q^{H_q(\delta, \hat{n}) \cdot \hat{n}}$, so there is always a choice of c_i outside these balls if we take $N = \lceil q^{(1 - H_q(\delta, \hat{n})) \cdot \hat{n}} \rceil$.

(2) We use the probabilistic method. Choose the N codewords randomly and independently from $\Sigma^{\hat{n}}$. The probability that there is a Hamming Ball of radius δ containing at least $L + 1$ codewords is at most

$$q^{\hat{n}} \cdot \binom{N}{L + 1} \cdot \left(\frac{q^{H_q(\delta, \hat{n}) \cdot \hat{n}}}{q^{\hat{n}}} \right)^{L + 1} \leq \left(\frac{N - 1}{q^{(1 - H_q(\delta, \hat{n}) - 1/(L + 1)) \cdot \hat{n}}} \right)^{L + 1},$$

which is less than 1 if we take $N = \lceil q^{(1 - H_q(\delta, \hat{n}) - 1/(L + 1)) \cdot \hat{n}} \rceil$. □

Note that while the rate bounds are essentially the same for achieving minimum distance δ and achieving list-decoding radius δ (as we take large list size), the bounds are incomparable because minimum distance δ only corresponds to unique decoding up to radius roughly $\delta/2$. The bound for list decoding is known to be tight up to the dependence on L (Problem 5.1), while the bound on minimum distance is not tight in general. Indeed, there are families of “algebraic-geometric” codes with constant alphabet size q , constant minimum distance $\delta > 0$, and constant rate $\rho > 0$ where $\rho > 1 - H_q(\delta, \hat{n})$ for sufficiently large \hat{n} . (Thus, this is a rare counterexample to the phenomenon “random is best”.) Identifying the best tradeoff between rate and minimum distance, even for binary codes, is a long-standing open problem in coding theory.

Open Problem 5.9. For each constant $\delta \in (0, 1)$, identify the largest $\rho > 0$ such that for every $\varepsilon > 0$, there exists an infinite family of codes $\mathcal{C}_{\hat{n}} \subset \{0, 1\}^{\hat{n}}$ of rate at least $\rho - \varepsilon$ and minimum distance at least δ .

Let’s look at some special cases of the parameters in Theorem 5.8. For binary codes ($q = 2$), we will be most interested in the case $\delta \rightarrow 1/2$, which corresponds to correcting the maximum possible

fraction of errors for binary codes. (No nontrivial decoding is possible for binary codes at distance greater than $1/2$, since a completely random received word will be at distance roughly $1/2$ with most codewords.) In this case, Proposition 5.7 tells us that the rate approaches $1 - H_2(1/2 - \varepsilon) = \Theta(\varepsilon^2)$, i.e. the blocklength is $\hat{n} = \Theta(n/\varepsilon^2)$ (for list size $L = \Theta(1/\varepsilon^2)$). For large alphabets q , Proposition 5.7 tells us that the rate approaches $1 - \delta$ as q grows. We will be most interested in the case $\delta = 1 - \varepsilon$ for small ε , so we can correct a $\delta = 1 - \varepsilon$ fraction of errors with a rate arbitrarily close to ε . For example, we can achieve rate of $\rho = .99\varepsilon$, a list size of $L = O(1/\varepsilon)$ and an alphabet of size $\text{poly}(1/\varepsilon)$. More generally, it is possible to achieve rate $\rho = (1 - \gamma)\varepsilon$ with an alphabet size of $q = (1/\varepsilon)^{O(1/\gamma)}$.

While we are primarily interested in list-decodable codes, minimum distance is often easier to bound. The following allows us to translate bounds on minimum distance into bounds on list-decodability.

Proposition 5.10 (Johnson Bound). (1) If \mathcal{C} has minimum distance $1 - \varepsilon$, then it is $(1 - O(\sqrt{\varepsilon}), O(1/\sqrt{\varepsilon}))$ -list-decodable.
(2) If a binary code \mathcal{C} has minimum distance $1/2 - \varepsilon$, then it is $(1/2 - O(\sqrt{\varepsilon}), O(1/\varepsilon))$ -list-decodable.

Proof. We prove Item 1, and leave Item 2 as an exercise in the next chapter (Problem 6.3). The proof is by inclusion-exclusion. Suppose for contradiction that there are codewords c_1, \dots, c_s at distance less than $1 - \varepsilon'$ from some $r \in \Sigma^{\hat{n}}$, for $\varepsilon' = 2\sqrt{\varepsilon}$ and $s = \lceil 2/\varepsilon' \rceil$. Then:

$$\begin{aligned}
1 &\geq \text{fraction of positions where } r \text{ agrees with some } c_i \\
&\geq \sum_i \text{agr}(r, c_i) - \sum_{1 \leq i < j \leq s} \text{agr}(c_i, c_j) \\
&> s\varepsilon' - \binom{s}{2} \cdot \varepsilon \\
&\geq 2 - 1 = 1
\end{aligned}$$

where the last inequality is by our setting of parameters. Contradiction. \square

Note the quadratic loss in the distance parameter. This means that optimal codes with respect to minimum distance are not necessarily optimal with respect to list decoding. Nevertheless, if we do not care about the exact tradeoff between the rate and the decoding radius, the above can yield codes where the decoding radius is as large as possible (approaching 1 for large alphabets and $1/2$ for binary alphabets).

5.1.3 Explicit Codes

As usual, most applications of error-correcting codes (in particular the original motivating one) require computationally efficient encoding and decoding. For now, we focus on only the efficiency of encoding.

Definition 5.11. A code $\text{Enc} : \{0, 1\}^n \rightarrow \Sigma^{\hat{n}}$ is (fully) *explicit* if given a message $m \in \{0, 1\}^n$ and an index $i \in \hat{n}$, the i 'th symbol of $\text{Enc}(m)$ can be computed in time $\text{poly}(n, \log \hat{n}, \log |\Sigma|)$.

The reason we talk about computing individual symbols of the codeword rather than the entire codeword is to have a meaningful definition even for codes where the blocklength \hat{n} is superpolynomial in the message length n . One can also consider weaker notions of explicitness, where we simply require that the entire codeword $\text{Enc}(m)$ can be computed in time $\text{poly}(\hat{n}, \log |\Sigma|)$.

The constructions of codes that we describe will involve arithmetic over finite fields. See Remark 3.25 for the complexity of constructing finite fields and carrying out arithmetic in such fields

In describing the explicit constructions below, it will be convenient to think of the codewords as functions $c : [\hat{n}] \rightarrow \Sigma$ rather than as strings in $\Sigma^{\hat{n}}$.

Construction 5.12 (Hadamard Code). For $n \in \mathbb{N}$, the (*binary*) *Hadamard code* of message length n is the binary code of blocklength $\hat{n} = 2^n$ consisting of all functions $c : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2$ that are *linear* (modulo 2).

Proposition 5.13. The Hadamard code:

- (1) is explicit with respect to the encoding function that takes a message $m \in \mathbb{Z}_2^n$ to the linear function c_m defined by $c_m(x) = \sum_i m_i x_i \pmod{2}$.
- (2) has minimum distance $1/2$, and
- (3) is $(1/2 - \varepsilon, O(1/\varepsilon^2))$ list-decodable for every $\varepsilon > 0$.

Proof. Explicitness is clear by inspection. The minimum distance follows from the fact that for every two distinct linear functions $c, c' : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2$, $\Pr_x[c(x) = c'(x)] = \Pr_x[(c - c')(x) = 0] = 1/2$. The list-decodability follows from the Johnson Bound. \square

The advantages of the Hadamard code are its small alphabet (binary) and optimal distance ($1/2$), but unfortunately its rate is exponentially small ($\rho = n/2^n$). By increasing both the field size and degree, we can obtain complementary properties

Construction 5.14 (Reed–Solomon Codes). For a prime power q and $d \in \mathbb{N}$, the q -ary *Reed–Solomon code* of degree d is the code of blocklength $\hat{n} = q$ and message length $n = (d + 1) \cdot \log q$ consisting of all polynomials $p : \mathbb{F}_q \rightarrow \mathbb{F}_q$ of degree at most d .

Proposition 5.15. The q -ary Reed–Solomon Code of degree d :

- (1) is explicit with respect to the encoding function that takes a vector of coefficients $m \in \mathbb{F}_q^{d+1}$ to the polynomial p_m defined by $p_m(x) = \sum_{i=0}^d m_i x^i$ (assuming we have a description of the field \mathbb{F}_q).
 - (2) has minimum distance $\delta = 1 - d/q$, and
 - (3) is $(1 - O(\sqrt{d/q}), O(\sqrt{q/d}))$ list-decodable.
-

Proof. Again explicitness follows by inspection. The minimum distance follows from the fact that two distinct polynomials of degree at most d agree in at most d points (else their difference would have more than d roots). The list-decodability follows from the Johnson Bound. \square

Note that by setting $q = O(d)$, Reed-Solomon codes simultaneously achieve constant rate and constant distance, the only disadvantage being that the alphabet is of nonconstant size (namely $q = \hat{n} > n$.)

Another useful setting of parameters for Reed-Solomon codes in complexity theory is $q = \text{poly}(d)$, which gives polynomial rate ($\hat{n} = \text{poly}(n)$) and distance tending to 1 polynomially fast ($\delta = 1 - 1/\text{poly}(n)$).

The following codes “interpolate” between Hadamard and Reed-Solomon codes, by allowing the number of variables, the degree, and field size all to vary.

Construction 5.16 (Reed–Muller Codes). For a prime power q and $d, m \in \mathbb{N}$, the q -ary *Reed–Muller code* of degree d and dimension m is the code of blocklength $\hat{n} = q^m$ and message length $n = \binom{m+d}{d} \cdot \log q$ consisting of all polynomials $p : \mathbb{F}_q^m \rightarrow \mathbb{F}_q$ of (total) degree at most d .

Proposition 5.17. The q -ary Reed–Muller Code of degree d and dimension m :

- (1) is explicit with respect to the encoding function that takes a vector of coefficients $v \in \mathbb{F}_q^{\binom{m+d}{d}}$ to the corresponding polynomial p_v .
- (2) has minimum distance $\delta \geq 1 - d/q$, and
- (3) is $(1 - O(\sqrt{d/q}), O(\sqrt{q/d}))$ list-decodable.

Proof. Same as for Reed–Solomon Codes, except we use the Schwartz-Zippel Lemma (Lemma 2.4) to deduce the minimum distance. \square

Note that Reed–Solomon Codes are simply Reed–Muller codes of dimension $m = 1$, and Hadamard codes are essentially Reed–Muller codes of degree $d = 1$ and alphabet size $q = 2$ (except that the Reed–Muller code also contains affine linear functions).

5.2 List-Decoding Algorithms

In this section, we will describe efficient list-decoding algorithms for the Reed–Solomon code and variants. It will be convenient to work with the following notation:

Definition 5.18. Let \mathcal{C} be a code with encoding function $\text{Enc} : \{1, \dots, N\} \rightarrow \Sigma^{\hat{n}}$. For $r \in \Sigma^{\hat{n}}$, define $\text{LIST}(r, \varepsilon) = \{m \in \Sigma^{\hat{n}} : \text{agr}(\text{Enc}(m), r) > \varepsilon\}$.

Then the task of $(1 - \varepsilon)$ list decoding (according to Definition 5.3) is equivalent to producing the elements of $\text{LIST}(r, \varepsilon)$ given $r \in \Sigma^{\hat{n}}$. In this section, we will see algorithms that do this in time polynomial in the bit-length of r , i.e. time $\text{poly}(\hat{n}, \log |\Sigma|)$.

5.2.1 Review of Algebra

The list-decoding algorithms will require some additional algebraic facts and notation:

- For every field \mathbb{F} , $\mathbb{F}[X_1, \dots, X_n]$ is the integral domain consisting of formal polynomials $Q(X_1, \dots, X_n)$ with coefficients in \mathbb{F} , where addition and multiplication of polynomials is defined in the usual way.
- A nonzero polynomial $Q(X_1, \dots, X_n)$ is *irreducible* if we cannot write $Q = RS$ where R, S are nonconstant polynomials. For a finite field \mathbb{F}_q of characteristic p and $d \in \mathbb{N}$, a univariate irreducible polynomial of degree d over \mathbb{F}_q can be found in deterministically in time $\text{poly}(p, \log q, d)$.
- $\mathbb{F}[X_1, \dots, X_n]$ is a *unique factorization domain*. That is, every nonzero polynomial Q can be factored as $Q = Q_1 Q_2 \cdots Q_m$, where each Q_i is irreducible and this factorization is unique up to reordering and multiplication by constants from \mathbb{F} . Given the description of a finite field \mathbb{F}_{p^k} and the polynomial Q , this factorization can be done probabilistically in time $\text{poly}(\log p, k, |Q|)$ and deterministically in time $\text{poly}(p, k, |Q|)$.
- For a nonzero polynomial $Q(Y, Z) \in \mathbb{F}[Y, Z]$ and $f(Y) \in \mathbb{F}[Y]$, if $Q(Y, f(Y)) = 0$, then $Z - f(Y)$ is one of the irreducible factors of $Q(Y, Z)$ (and thus $f(Y)$ can be found in polynomial time). This is analogous to the fact that if $c \in \mathbb{Z}$ is a root of an integer polynomial $Q(Z)$, then $Z - c$ is one of the factors of Q (and can be proven in the same way, by long division).

5.2.2 List-Decoding Reed-Solomon Codes

Theorem 5.19. ~~There is a polynomial-time $(1 - \varepsilon)$ list-decoding algorithm for the q -ary Reed-Solomon code of degree d , for $\varepsilon = 2\sqrt{d/q}$. That is, given a function $r : \mathbb{F}_q \rightarrow \mathbb{F}_q$ and $d \in \mathbb{N}$, all polynomials of degree at most d that agree with r on more than $\varepsilon q = 2\sqrt{dq}$ inputs can be found in polynomial time.~~

In fact the constant of 2 can be improved to 1, matching the combinatorial list-decoding radius for Reed-Solomon codes given by an optimized form of the Johnson Bound. See Problem 5.6.

Proof. We are given a received word $r : \mathbb{F}_q \rightarrow \mathbb{F}_q$, and want to find all elements of $\text{LIST}(r, \varepsilon)$ for $\varepsilon = 2\sqrt{d/q}$.

Step 1: Find a low-degree Q “explaining” r . Specifically, $Q(Y, Z)$ will be a nonzero bivariate polynomial of degree at most d_Y in its first variable Y and d_Z in its second variable, with the property that $Q(y, r(y)) = 0$ for all $y \in \mathbb{F}_q$. Each such y imposes a linear constraint on the $(d_Y+1)(d_Z+1)$ coefficients of Q . Thus, this system has a nonzero solution provided $(d_Y+1)(d_Z+1) > q$, and it can be found in polynomial time by linear algebra (over \mathbb{F}_q).

Step 2: Argue that each $f(Y) \in \text{LIST}(r, \varepsilon)$ is a “root” of Q . Specifically, it will be the case that $Q(Y, f(Y)) = 0$ for each $f \in \text{LIST}(r, \varepsilon)$. The reason is that $Q(Y, f(Y))$ is a univariate polynomial of degree at most $d_Y + d \cdot d_Z$, and has more than εq zeroes (one for each place that f and r agree). Thus, we can conclude $Q(Y, f(Y)) = 0$ provided $\varepsilon q \geq d_Y + d \cdot d_Z$. Then we can

enumerate all of the elements of $\text{LIST}(r, \varepsilon)$ by factoring $Q(Y, Z)$ and taking all the factors of the form $Z - f(Y)$.

For this algorithm to work, the two conditions we need to satisfy are

$$(d_Y + 1)(d_Z + 1) > q,$$

and

$$\varepsilon q \geq d_Y + d \cdot d_Z.$$

These conditions can be satisfied by setting $d_Y = \lfloor \varepsilon q / 2 \rfloor$, $d_Z = \lfloor \varepsilon q / (2d) \rfloor$, and $\varepsilon = 2\sqrt{d/q}$. \square

Note that the rate of Reed-Solomon codes is $\rho = (d + 1)/q = \Theta(\varepsilon^2)$. The alphabet size is $q = \tilde{\Omega}(n/\rho) = \tilde{\Omega}(n/\varepsilon^2)$. In contrast, the random codes of Theorem 5.8 achieve $\rho \approx \varepsilon$ and $q = \text{poly}(1/\varepsilon)$. It is not known whether the $\varepsilon = \sqrt{d/q}$ bound on the list-decodability of Reed-Solomon codes can be improved, even with inefficient decoding.

Open Problem 5.20. Do there exist constants $\varepsilon, \rho \in (0, 1)$ with $\varepsilon < \sqrt{\rho}$ and an infinite sequence of prime powers q such that the q -ary Reed-Solomon code of degree $d = \lfloor \rho q \rfloor$ is $(1 - \varepsilon, \text{poly}(q))$ -list-decodable?

5.2.3 Parvaresh-Vardy Codes

Our aim is to improve the rate-distance tradeoff to $\rho = \tilde{\Theta}(\varepsilon)$. Intuitively, the power of the Reed-Solomon list-decoding algorithm comes from the fact that we can interpolate the q points $(y, r(y))$ of the received word using a *bivariate* polynomial Q of degree roughly \sqrt{q} in each variable (think of $d = O(1)$ for now). If we could use m variables instead of 2, then the degrees would only have to be around $q^{1/m}$.

First attempt: Replace Step 1 with finding an $(m+1)$ -variate polynomial $Q(Y, Z_1, \dots, Z_m)$ of degree d_Y in Y and d_Z in each Z_i such that $Q(y, r(y), r(y), \dots, r(y)) = 0$ for every $y \in \mathbb{F}_q$. Then, we will be able to choose a nonzero polynomial Q of degree roughly $q^{1/m}$ such that $Q(Y, f(Y), \dots, f(Y)) = 0$ for every $f \in \text{LIST}(r, \varepsilon)$, and hence it follows that $Z - f(Y)$ divides the bivariate polynomial $Q^*(Y, Z) = Q(Y, Z, \dots, Z)$. Unfortunately, Q^* might be the zero polynomial even if Q is nonzero.

Second attempt: Replace Step 1 with finding an $(m + 1)$ -variate polynomial $Q(Y, Z_1, \dots, Z_m)$ of degree d_Y in Y and $d_Z = h - 1$ in each Z_i such that $Q(y, r(y), r(y)^h, r(y)^{h^2}, \dots, r(y)^{h^{m-1}}) = 0$ for every $y \in \mathbb{F}_q$. Then, it follows that $Q^*(Y, Z) = Q(Y, Z, Z^h, \dots, Z^{h^{m-1}})$ is nonzero if Q is nonzero because every monomial in Q with individual degrees at most $h - 1$ in Z_1, \dots, Z_m gets mapped to a different power of Z . However, here the difficulty is that the degree of $Q^*(Y, f(Y))$ is too high (roughly $d^* = d_Y + d \cdot h^m > d_Z^m$) for us to satisfy the constraint $\varepsilon q \geq d^*$.

Parvaresh-Vardy codes get the best of both worlds by providing more information with each symbol — not just the evaluation of f at each point, but the evaluation of $m - 1$ other polynomials f_1, \dots, f_{m-1} , each of which is still of degree d (as is good for arguing that

$Q(Y, f(Y), f_1(Y), \dots, f_{m-1}(Y)) = 0$, but can be viewed as raising f to successive powers of h for the purposes of ensuring that Q^* is nonzero.

To introduce this idea, we need some additional algebra.

- For univariate polynomials $f(Y)$ and $E(Y)$, we define $f(Y) \bmod E(Y)$ to be the remainder when f is divided by E . If $E(Y)$ is of degree k , then $f(Y) \bmod E(Y)$ is of degree at most $k - 1$.
- The ring $\mathbb{F}[Y]/E(Y)$ consists of all polynomials of degree at most $k - 1$ with arithmetic modulo $E(Y)$ (analogous to \mathbb{Z}_n consisting of integers smaller than n with arithmetic modulo n). If E is irreducible, then $\mathbb{F}[Y]/E(Y)$ is a field (analogous to \mathbb{Z}_p being a field when p is prime). Indeed, this is how the finite field of size p^k is constructed: take $\mathbb{F} = \mathbb{Z}_p$ and $E(Y)$ to be an irreducible polynomial of degree k over \mathbb{Z}_p , and then $\mathbb{F}[Y]/E(Y)$ is the (unique) field of size p^k .
- A multivariate polynomial $Q(Y, Z_1, \dots, Z_m)$ can be reduced modulo $E(Y)$ by writing it as a polynomial in variables Z_1, \dots, Z_m with coefficients in $\mathbb{F}[Y]$ and then reducing each coefficient modulo $E(Y)$. After reducing Q modulo E , we think of Q as a polynomial in variables Z_1, \dots, Z_m with coefficients in the field $\mathbb{F}[Y]/E(Y)$.

Construction 5.21 (Parvaresh–Vardy Codes). For a prime power q , integers $m, d, h \in \mathbb{N}$, and an irreducible polynomial $E(Y)$ of degree larger than d , the q -ary *Parvaresh–Vardy Code* of degree d , power h , redundancy m , and irreducible E is defined as follows:

- The alphabet is $\Sigma = \mathbb{F}_q^m$.
- The blocklength is $\hat{n} = q$.
- The message space is \mathbb{F}_q^{d+1} , where we view each message as representing a polynomial $f(Y)$ of degree at most d over \mathbb{F}_q .
- For $y \in \mathbb{F}_q$, the y 'th symbol of the $\text{Enc}(f)$ is

$$[f_0(y), f_1(y), \dots, f_{m-1}(y)],$$

where $f_i(Y) = f(Y)^{h^i} \bmod E(Y)$.

Theorem 5.22. For every prime power q , integer $0 \leq d < q$, and irreducible polynomial E of degree $d + 1$, the q -ary Parvaresh–Vardy code of degree d , redundancy $m = \lceil \log(q/d) \rceil$, power $h = 2$, and irreducible E has rate $\rho = \tilde{\Omega}(d/q)$ and can be list-decoded in polynomial time up to distance $\delta = 1 - \tilde{O}(d/q)$.

Proof. We are given a received word $r : \mathbb{F}_q \rightarrow \mathbb{F}_q^m$, and want to find all elements of $\text{LIST}(r, \varepsilon)$, for some $\varepsilon = \tilde{O}(d/q)$.

Step 1: Find a low-degree Q “explaining” r . We find a nonzero polynomial $Q(Y, Z_0, \dots, Z_{m-1})$ of degree at most d_Y in its first variable Y and at most $h - 1$ in each of the remaining variables, and satisfying $Q(y, r(y)) = 0$ for all $y \in \mathbb{F}_q$.

Such a Q exists and can be found by linear algebra in time $\text{poly}(q, m)$ provided that:

$$d_Y \cdot h^m > q. \quad (5.1)$$

Moreover, we may assume that Q is not divisible by $E(Y)$. If it is, we can divide out all the factors of $E(Y)$, which will not affect the conditions $Q(y, r(y)) = 0$ since E has no roots (being irreducible).

Step 2: Argue that each $f(Y) \in \text{LIST}(r, \varepsilon)$ is a “root” of a related univariate polynomial Q^* . First, we argue as before that for $f \in \text{LIST}(r, \varepsilon)$, we have

$$Q(Y, f_0(Y), \dots, f_{m-1}(Y)) = 0. \quad (5.2)$$

Since each f_i has degree at most $\deg(E) - 1 = d$, this will be ensured provided

$$\varepsilon q \geq d_Y + (h - 1) \cdot d \cdot m. \quad (5.3)$$

Once we have this, we can reduce both sides of Equation (5.2) modulo $E(Y)$ and deduce

$$\begin{aligned} 0 &= Q(Y, f_0(Y), f_1(Y), \dots, f_{m-1}(Y)) \bmod E(Y) \\ &= Q(Y, f(Y), f(Y)^h, \dots, f(Y)^{h^{m-1}}) \bmod E(Y) \end{aligned}$$

Thus, if we define the univariate polynomial

$$Q^*(Z) = Q(Y, Z, Z^h, \dots, Z^{h^{m-1}}) \bmod E(Y),$$

then $f(Y)$ is a root of Q^* over the field $\mathbb{F}_q[Y]/E(Y)$.

Observe that Q^* is nonzero because Q is not divisible by $E(Y)$ and has degree at most $h - 1$ in each Z_i . Thus, we can find all elements of $\text{LIST}(r, \varepsilon)$ by factoring $Q^*(Z)$. Constructing $Q^*(Z)$ and factoring it can be done in time $\text{poly}(q, d, h^m) = \text{poly}(q)$.

For this algorithm to work, we need to satisfy Conditions (5.1) and (5.3). We can satisfy Condition (5.3) by setting $d_Y = \lceil \varepsilon q - dhm \rceil$, in which case Condition (5.1) is satisfied for

$$\varepsilon = \frac{1}{h^m} + \frac{dhm}{q} = \tilde{O}(d/q), \quad (5.4)$$

for $h = 2$ and $m = \lceil \log(q/d) \rceil$. Observing that the rate is $\rho = d/(mq) = \tilde{\Omega}(d/q)$, this completes the proof of the theorem. \square

Note that the obstacles to obtaining a tradeoff $\rho \approx \varepsilon$ are the factor of m in expression for the rate $\rho = d/(mq)$ and the factor of hm in Equation (5.4) for ε . We remedy these in the next section.

5.2.4 Folded Reed–Solomon Codes

In this section, we show how to obtain an optimal rate-distance tradeoff, where the rate ρ is arbitrarily close to the agreement parameter ε .

Consider the Parvaresh–Vardy construction with polynomial $E(Y) = Y^{q-1} - \gamma$, where γ is a generator of the multiplicative group \mathbb{F}_q^* ; it can be shown that $E(Y)$ is irreducible. (That is, $\{\gamma, \gamma^2, \dots, \gamma^{q-1}\} = \mathbb{F}_q \setminus \{0\}$.) Then it turns out that $f^q(Y) \bmod E(Y) = f(Y^q) \bmod E(Y) = f(\gamma Y)$.

So, we set $h = q$ and the degree of $f_i(Y) = f^{h^i}(Y) \bmod E(Y) = f(\gamma^i Y)$ is d even though E has degree $q - 1$. For each nonzero element y of \mathbb{F}_q , the y 'th symbol of the PV encoding of $f(Y)$ is then

$$[f(y), f(\gamma y), \dots, f(\gamma^{m-1}y)] = [f(\gamma^j), f(\gamma^{j+1}), \dots, f(\gamma^{j+m-1})], \quad (5.5)$$

where we write $y = \gamma^j$.

Thus, the symbols of the PV encoding have a lot of overlap. For example, the γ^j 'th symbol and the γ^{j+1} 'th symbol share all but one component. Intuitively, this means that we should only have to send a $1/m$ fraction of the symbols of the codeword, saving us a factor of m in the rate. (The other symbols can be automatically filled in by the receiver.) Thus, the rate becomes $\rho = d/q$, just like in Reed–Solomon codes.

More formally, we use the following codes.

Construction 5.23 (Folded Reed–Solomon Codes). For a prime power q , a generator γ of \mathbb{F}_q^* , integers $m, d \in \mathbb{N}$, the q -ary *folded Reed–Solomon code* of degree d , folding parameter m , and generator γ is defined as follows:

- The alphabet is $\Sigma = \mathbb{F}_q^m$.
- The blocklength is $\hat{n} = \lfloor (q - 1)/m \rfloor$.
- The message space is \mathbb{F}_q^{d+1} , where we view each message as representing a polynomial $f(Y)$ of degree at most d over \mathbb{F}_q .
- For $k \in [\hat{n}]$, the k 'th symbol of $\text{Enc}(f)$ is

$$[f(\gamma^{(k-1)m}), f(\gamma^{(k-1)m+1}), \dots, f(\gamma^{k \cdot m-1})].$$

We now show that these codes can be efficiently list-decoded at distance arbitrarily close to $1 - \rho$, where $\rho = d/q$ is the rate.

Theorem 5.24. For every prime power q , integers $0 \leq d, m \leq q$, and generator γ of \mathbb{F}_q^* , the q -ary folded Reed–Solomon code of degree d , folding parameter m , and generator γ has rate at least $\rho = d/q$ and can be list-decoded in time $q^{O(m)}$ up to distance $\delta = 1 - d/q - O(\sqrt{1/m})$.

Note that this theorem is most interesting (and improves on the Reed–Solomon decoding of Theorem 5.19) when m is larger than q/d , in contrast to the setting of parameters in Theorem 5.22, where m is logarithmic in q . We can afford the larger setting of m because now the rate ρ does not depend on m , and it is this change in parameters that enables us to save us the factor of hm in Equation (5.4). However, the running time, list size, and alphabet size do grow exponentially with m .

Proof. We are given a received word $r : [\hat{n}] \rightarrow \mathbb{F}_q^m$, and want to find all elements of $\text{LIST}(r, \varepsilon)$ for an appropriate choice of $\varepsilon = d/q + O(\sqrt{1/m})$.

Step 0: Unpack to a PV received word. From r , we will obtain a received word $r' : \mathbb{F}_q \rightarrow \mathbb{F}_q^{m'}$ for the Parvaresh–Vardy code of degree d , power $h = q$, redundancy $m' = \lfloor \sqrt{m-1} \rfloor$, and irreducible

$E(Y) = Y^{q-1} - \gamma$. Specifically, following Equation (5.5), each symbol of r yields $m - m'$ symbols of r' . If r agrees with the FRS encoding of f in more than $\varepsilon \hat{n}$ positions, then r' will agree with the PV encoding of f in more than $\varepsilon \hat{n} \cdot (m - m')$ positions. We have

$$\varepsilon \hat{n} \cdot (m - m') \geq \varepsilon \cdot ((q - 1)/m - 1) \cdot (m - m') = (\varepsilon - O(1/\sqrt{m})) \cdot q.$$

Thus our task is now to find $\text{LIST}(r', \varepsilon')$ for $\varepsilon' = \varepsilon - O(1/\sqrt{m})$, which we do in a similar manner to the decoding algorithm for Parvaresh–Vardy codes.

Step 1: Find a low-degree Q “explaining” r' . We find a nonzero polynomial $Q(Y, Z_0, \dots, Z_{m'-1})$ satisfying $Q(y, r'(y)) = 0$ for all $y \in \mathbb{F}_q$ and such that Q has degree at most d_Y in its first variable Y and *total degree* at most $d_Z = 1$ in the Z_i variables. That is, Q only has monomials of the form $Y^j Z_i$ for $j \leq d_Y$. The number of these monomials is $(d_Y + 1) \cdot m'$, so we can find such a Q in time $\text{poly}(q, m') = \text{poly}(q, m)$ provided:

$$(d_Y + 1) \cdot m' > q. \tag{5.6}$$

As before, we may assume that Q is not divisible by $E(Y)$. (Note that using total degree 1 instead of individual degrees $h - 1$ in the Z_i 's requires an exponentially larger setting of m' compared to the setting of m needed for Equation (5.1). However, this will provide a significant savings below. In general, it is best to consider a combination of the two constraints, requiring that the total degree in the Z_i 's is at most some d_Z and that the individual degrees are all at most $h - 1$.)

Step 2: Argue that each $f(Y) \in \text{LIST}(r', \varepsilon')$ is a “root” of a related univariate polynomial Q^* . First, we argue as before that for $f \in \text{LIST}(r', \varepsilon')$, we have

$$Q(Y, f_0(Y), \dots, f_{m'-1}(Y)) = 0, \tag{5.7}$$

where $f_i(Y) = f^{h^i}(Y) \bmod E(Y) = f(\gamma^i Y)$. Since each f_i has degree at most d , this will be ensured provided

$$\varepsilon' q \geq d_Y + d. \tag{5.8}$$

Note the savings of the factor $(h - 1) \cdot m$ as compared to Inequality (5.3); this is because we chose Q to be of total degree 1 in the Z_i 's instead of having individual degree 1.

Now, as in the Parvaresh–Vardy decoding, we can reduce both sides of Equation (5.7) modulo $E(Y)$ and deduce

$$\begin{aligned} 0 &= Q(Y, f_0(Y), f_1(Y), \dots, f_{m'-1}(Y)) \bmod E(Y) \\ &= Q(Y, f(Y), f(Y)^h, \dots, f(Y)^{h^{m-1}}) \bmod E(Y). \end{aligned}$$

Thus, if we define the univariate polynomial

$$Q^*(Z) = Q(Y, Z, Z^h, \dots, Z^{h^{m-1}}) \bmod E(Y),$$

then $f(Y)$ is a root of Q^* over the field $\mathbb{F}_q[Y]/E(Y)$.

Observe that Q^* is nonzero because Q is not divisible by $E(Y)$ and has degree at most $h - 1$ in each Z_i . Thus, we can find all elements of $\text{LIST}(r', \varepsilon')$ by factoring $Q^*(Z)$. Constructing and factoring $Q^*(Z)$ can be done in time $\text{poly}(q, h^m) = q^{O(m)}$.

For this algorithm to work, we need to satisfy Conditions (5.6) and (5.8). We can satisfy Condition (5.6) by setting $d_Y = \lfloor q/m' \rfloor$, in which case Condition (5.8) is satisfied for

$$\varepsilon' \geq 1/m' + d/q.$$

Recalling that $\varepsilon' = \varepsilon - O(1/\sqrt{m})$ and $m' = \lfloor \sqrt{m-1} \rfloor$, we can take $\varepsilon = d/q + O(1/\sqrt{m})$. \square

Setting parameters appropriately gives:

Theorem 5.25. The following holds for all *constants* $\varepsilon, \gamma > 0$. For every $n \in \mathbb{N}$, there is an explicit code of message length n and rate $\rho = \varepsilon - \gamma$ that can be list-decoded in polynomial time from distance $1 - \varepsilon$, with alphabet size $q = \text{poly}(n)$.

The polynomials in the running time, alphabet size, and list size depend exponentially on γ ; specifically they are of the form $n^{O(1/\gamma^2)}$. Nonconstructively, it is possible to have alphabet size $(1/\varepsilon)^{O(\varepsilon/\gamma)}$ and list size $O(1/\gamma)$ (see discussion after Theorem 5.8), and one could hope for running time that is a fixed polynomial in n , with an exponent that is independent of γ . Recent work has achieved these parameters, albeit with a randomized construction of codes; it remains open to have a fully explicit and deterministic construction. However, for a fixed constant-sized alphabet, e.g. $q = 2$, it is still not known how to achieve list-decoding capacity.

Open Problem 5.26. For any desired constants $\rho, \delta > 0$ such that $\rho > 1 - H_2(\delta)$, construct an explicit family of codes $\text{Enc}_n : \{0, 1\}^n \rightarrow \{0, 1\}^{\hat{n}}$ that have rate at least ρ and are δ -list-decodable in polynomial time.

5.3 List-decoding views of samplers and expanders

In this section, we show how averaging samplers and expander graphs can be understood as variants of list-decodable codes. The general list decoding framework we use to describe these connections will also allow us to capture the other pseudorandom objects we study in this survey.

We begin with the syntactic correspondence between the three objects, so we can view each as a function $\Gamma : [N] \times [D] \rightarrow [M]$:

Construction 5.27 (Syntactic Correspondence of Expanders, Samplers, and Codes).

- (1) Given a bipartite multigraph G with N left vertices, M right vertices, and left-degree D , we let $\Gamma : [N] \times [D] \rightarrow [M]$ be its neighbor function, i.e. $\Gamma(x, y)$ is the y 'th neighbor of x in G .
- (2) Given a sampler $\text{Samp} : [N] \rightarrow [M]^D$, we define $\Gamma : [N] \times [D] \rightarrow [M]$ by

$$\Gamma(x, y) = \text{Samp}(x)_y.$$

- (3) Given a code $\text{Enc} : [N] \rightarrow [q]^D$, we set $M = q \cdot D$, associate the elements of $[q] \times [D]$ with $[M]$, and define $\Gamma : [N] \times [D] \rightarrow [M]$ by

$$\Gamma(x, y) = (y, \text{Enc}(x)_y).$$

The correspondence between expanders and samplers is identical to the one from Problem 4.7, which shows that the vertex expansion of G is equivalent to the hitting sampler properties of Samp.

Note that codes correspond to expanders and samplers where the first component of a neighbor/sample equals the edge-label/sample-number. Conversely, any such expander/sampler can be viewed as a code. Many constructions of expanders and samplers have or can be easily modified to have this property. This syntactic constraint turns out to be quite natural in the setting of samplers. It corresponds to the case where we have D functions $f_1, \dots, f_D : [M] \rightarrow [0, 1]$, we are interested in estimating the total average $(1/D) \sum_i \mu(f_i)$, but can only evaluate each f_i on a single sample.

We now consider the following generalized notion of list decoding.

Definition 5.28. For a function $\Gamma : [N] \times [D] \rightarrow [M]$, a set $T \subset [M]$, and $\varepsilon \in [0, 1]$, we define

$$\begin{aligned} \text{LIST}_\Gamma(T, \varepsilon) &= \{x : \Pr_y[\Gamma(x, y) \in T] > \varepsilon\} \text{ and} \\ \text{LIST}_\Gamma(T, 1) &= \{x : \forall y \Gamma(x, y) \in T\}. \end{aligned}$$

More generally, for a function $f : [M] \rightarrow [0, 1]$, we define

$$\text{LIST}_\Gamma(f, \varepsilon) = \{x : \mathbb{E}_y[f(\Gamma(x, y))] > \varepsilon\}.$$

We can formulate the list-decoding property of Enc, the averaging sampler property of Samp, and the vertex expansion property of Γ in this language as follows:

Proposition 5.29. Let Enc and Γ be as in Construction 5.27. Then Enc is $(1 - 1/q - \varepsilon, K)$ list-decodable iff for every $r \in [M]^D$, we have

$$|\text{LIST}_\Gamma(T_r, 1/q + \varepsilon)| \leq K,$$

where $T_r = \{(y, r_y) : y \in [D]\}$.

Proof. Observe that, for each $x \in [N]$ and $r \in [M]^D$,

$$\begin{aligned} \Pr_y[\Gamma(x, y) \in T_r] &= \Pr_y[(y, \text{Enc}(x)_y) \in T_r] \\ &= \Pr_y[\text{Enc}(x)_y = r_y] \\ &= \text{agr}(\text{Enc}(x), r). \end{aligned}$$

Thus $|\text{LIST}(T_r, 1/q + \varepsilon)| \leq K$ if and only if there are at most K messages x whose encodings have agreement greater than $1/q + \varepsilon$ with r , which is the same as being $(1 - 1/q - \varepsilon, K)$ list-decodable. \square

Proposition 5.30. Let Samp and Γ be as in Construction 5.27. Then

(1) Samp is a (δ, ε) averaging sampler iff for every function $f : [M] \rightarrow [0, 1]$, we have

$$|\text{LIST}_\Gamma(f, \mu(f) + \varepsilon)| \leq \delta N.$$

(2) Samp is a (δ, ε) boolean averaging sampler iff for every set $T \subset [M]$, we have

$$|\text{LIST}_\Gamma(T, \mu(T) + \varepsilon)| \leq \delta N.$$

Proof. Observe that $x \in \text{LIST}_\Gamma(f, \mu(f) + \varepsilon)$ if and only if x is a “bad” set of coin tosses for Samp — namely $(1/D) \cdot \sum_{i=1}^D f(z_i) > \mu(f) + \varepsilon$, where $(z_1, \dots, z_D) = \text{Samp}(x)$. Thus Samp errs with probability at most δ over $x \stackrel{\text{R}}{\leftarrow} U_{[N]}$ if and only if $|\text{LIST}_\Gamma(f, \mu(f) + \varepsilon)| \leq \delta N$. The boolean averaging sampler case follows by viewing boolean functions f as characteristic functions of sets T . \square

Noting that the sets T_r in Proposition 5.29 have density $\mu(T_r) = 1/q$, we see that the averaging-sampler property implies the standard list-decoding property:

Corollary 5.31. If Samp is a (δ, ε) boolean averaging sampler of the form $\text{Samp}(x)_y = (y, \text{Enc}(x)_y)$, then Enc is $(1 - 1/q - \varepsilon, \delta N)$ list-decodable.

Note, however, that the typical settings of parameters of samplers and list-decodable codes are very different. With codes, we want the alphabet size q to be as small as possible (e.g. $q = 2$) and the blocklength D to be linear or polynomial in the message length $n = \log N$, so $M = qD$ is also linear or polynomial in $n = \log N$. In contrast, we usually are interested in samplers for functions on exponentially large domains (e.g. $M = 2^{\Omega(n)}$).

In Chapter 6, we will see a converse to Corollary 5.31 when the alphabet size is small: if Enc is $(1 - 1/q - \varepsilon, \delta N)$, list-decodable, then Samp is an $(\delta/\varepsilon, q \cdot \varepsilon)$ averaging sampler.

For expanders, it will be convenient to state the list-decoding property in terms of the following variant of vertex expansion, where we only require that sets of size exactly K expand:

Definition 5.32. For $K \in \mathbb{N}$, a bipartite multigraph G is an $(= K, A)$ *vertex expander* if all sets S consisting of K left-vertices, the neighborhood $N(S)$ is of size at least $A \cdot K$.

Thus, G is a (K, A) vertex expander in the sense of Definition 4.3 iff G is an $(= K', A)$ vertex expander for all positive integers $K' \leq K$.

Proposition 5.33. For $K \in \mathbb{N}$, $\Gamma : [N] \times [D] \rightarrow [M]$ is an $(= K, A)$ vertex expander iff for every set $T \subset [D] \times [M]$ such that $|T| < KA$, we have:

$$|\text{LIST}_\Gamma(T, 1)| < K.$$

Proof.

$$\begin{aligned} \Gamma \text{ not an } (= K, A) \text{ expander} &\Leftrightarrow \exists S \subset [N] \text{ s.t. } |S| = K \text{ and } |N(S)| < KA \\ &\Leftrightarrow \exists S \subset [N] \text{ s.t. } |S| \geq K \text{ and } |N(S)| < KA \\ &\Leftrightarrow \exists T \subset [M] \text{ s.t. } |\text{LIST}(T, 1)| \geq K \text{ and } |T| < KA, \end{aligned}$$

where the last equivalence follows because if $T = N(S)$, then $S \subset \text{LIST}(T, 1)$, and conversely if $S = \text{LIST}(T, 1)$ then $N(S) \subset T$. \square

On one hand, this list-decoding property seems easier to establish than the ones for codes and samplers because we look at $\text{LIST}(T, 1)$ instead of $\text{LIST}(T, \mu(T) + \varepsilon)$. On the other hand, to get expansion (i.e. $A > 1$), we require a very tight relationship between $|T|$ and $|\text{LIST}(T, 1)|$. In the setting of codes or samplers, we would not care much about a factor of 2 loss in $|\text{LIST}(T)|$, as this just corresponds to a factor of 2 in list size or error probability. But here it corresponds to a factor of 2 loss in expansion, which can be quite significant. In particular, we cannot afford it if we are trying to get $A = (1 - \varepsilon) \cdot D$, as we will be in the next section.

5.4 Expanders from Parvaresh–Vardy Codes

Despite the substantial differences in the standard settings of parameters between codes, samplers, and expanders, it can be very useful to translate ideas and techniques from one object to the other using the connections described in the previous section. In particular, in this section we will see how to build graphs with extremely good vertex expansion ($A = (1 - \varepsilon)D$) from Parvaresh–Vardy codes.

Consider the bipartite multigraph obtained from the Parvaresh–Vardy codes (Construction 5.21) via the correspondence of Construction 5.27. That is, we define a neighbor function $\Gamma : \mathbb{F}_q^n \times \mathbb{F}_q \rightarrow \mathbb{F}_q \times \mathbb{F}_q^m$ by

$$\Gamma(f, y) = [y, f_0(y), f_1(y), \dots, f_{m-1}(y)], \quad (5.9)$$

where $f(Y)$ is a polynomial of degree at most $n - 1$ over \mathbb{F}_q , and we define $f_i(Y) = f(Y)^{h^i} \bmod E(Y)$, where E is a fixed irreducible polynomial of degree n over \mathbb{F}_q . (Note that we are using $n - 1$ instead of d to denote degree of f .)

Theorem 5.34. Let $\Gamma : \mathbb{F}_q^n \times \mathbb{F}_q \rightarrow \mathbb{F}_q^{m+1}$ be the neighbor function of the bipartite multigraph corresponding to the q -ary Parvaresh–Vardy code of degree $d = n - 1$, power h , and redundancy m via Construction 5.27. Then Γ is a (K_{max}, A) expander for $K_{max} = h^m$ and $A = q - nhm$.

Proof. Let K be any integer less than or equal to $K_{max} = h^m$, and let $A = q - nmh$. By Proposition 5.33, it suffices to show that for every set $T \subset \mathbb{F}_q^{m+1}$ of size at most $AK - 1$, we have $|\text{LIST}(T)| \leq K - 1$.

We begin by doing the proof for $K = K_{max} = h^m$, and later describe the modifications to handle smaller values of K . The proof goes along the same lines as the list-decoding algorithm for the Parvaresh–Vardy codes from Section 5.2.3.

Step 1: Find a low-degree Q vanishing on T . We find a nonzero polynomial $Q(Y, Z_0, \dots, Z_{m-1})$ of degree at most $d_Y = A - 1$ in its first variable Y and at most $h - 1$ in each of the remaining variables such that $Q(z) = 0$ for all $z \in T$. (Compare this to $Q(y, r(y)) = 0$ for all $y \in \mathbb{F}_q$ in the list-decoding algorithm, which corresponds to taking $T = T_r$.)

This is possible because

$$A \cdot h^m = AK > |T|.$$

Moreover, we may assume that Q is not divisible by $E(Y)$. If it is, we can divide out all the factors of $E(Y)$, which will not affect the conditions $Q(z) = 0$ since E has no roots (being irreducible).

Step 2: Argue that each $f(Y) \in \text{LIST}(T, 1)$ is a “root” of a related univariate polynomial Q^* . First, we argue as in the list-decoding algorithm that if $f \in \text{LIST}(T, 1)$, we have

$$Q(Y, f_0(Y), \dots, f_{m-1}(Y)) = 0.$$

This is ensured because

$$q > A - 1 + nmh.$$

(In the list-decoding algorithm, the left-hand side of this inequality was εq , since we were bounding $|\text{LIST}(T_r, \varepsilon)|$.)

Once we have this, we can reduce both sides modulo $E(Y)$ and deduce

$$\begin{aligned} 0 &= Q(Y, f_0(Y), f_1(Y), \dots, f_{m-1}(Y)) \bmod E(Y) \\ &= Q(Y, f(Y), f(Y)^{h^2}, \dots, f(Y)^{h^{m-1}}) \bmod E(Y) \end{aligned}$$

Thus, if we define the univariate polynomial

$$Q^*(Z) = Q(Y, Z, Z^h, \dots, Z^{h^{m-1}}) \bmod E(Y),$$

then $f(Y)$ is a root of Q^* over the field $\mathbb{F}_q[Y]/E(Y)$.

Observe that Q^* is nonzero because Q is not divisible by $E(Y)$ and has degree at most $h - 1$ in each Z_i . Thus,

$$|\text{LIST}(T, 1)| \leq \deg(Q^*) \leq h - 1 + (h - 1) \cdot h + (h - 1) \cdot h^2 + \dots + (h - 1) \cdot h^{m-1} = K - 1.$$

(Compare this to the list-decoding algorithm, where our primary goal was to efficiently enumerate the elements of $\text{LIST}(T, \varepsilon)$, as opposed to bound its size.)

Handling smaller values of K . We further restrict $Q(Y, Z_1, \dots, Z_m)$ to only have nonzero coefficients on monomials of the form $Y^i \text{Mon}_j(Z_1, \dots, Z_m)$ for $0 \leq i \leq A - 1$ and $0 \leq j \leq K - 1 \leq h^m - 1$, where $\text{Mon}_j(Z_1, \dots, Z_m) = Z_1^{j_0} \dots Z_m^{j_{m-1}}$ and $j = j_0 + j_1 h + \dots + j_{m-1} h^{m-1}$ is the base- h representation of j . Note that this gives us $AK > |T|$ monomials, so Step 1 is possible. Moreover $M_j(Z, Z^h, Z^{h^2}, \dots, Z^{h^{m-1}}) = Z^j$, so the degree of Q^* is at most $K - 1$, and we get the desired list-size bound in Step 3. \square

Setting parameters, we have:

Theorem 5.35. For every constant $\alpha > 0$, every $N \in \mathbb{N}$, $K \leq N$, and $\varepsilon > 0$, there is an explicit $(K, (1 - \varepsilon)D)$ expander with N left-vertices, M right-vertices, left-degree $D = O((\log N)(\log K)/\varepsilon)^{1+1/\alpha}$ and $M \leq D^2 \cdot K^{1+\alpha}$. Moreover, D is a power of 2.

Proof. Let $n = \log N$ and $k = \log K_{\max}$. Let $h = \lceil (2nk/\varepsilon)^{1/\alpha} \rceil$ and let q be the power of 2 in the interval $(h^{1+\alpha}/2, h^{1+\alpha}]$.

Set $m = \lceil (\log K_{\max})/(\log h) \rceil$, so that $h^{m-1} \leq K_{\max} \leq h^m$. Then, by Theorem 5.34, the graph $\Gamma : \mathbb{F}_q^n \times \mathbb{F}_q \rightarrow \mathbb{F}_q^{m+1}$ defined in (5.9) is an (h^m, A) expander for $A = q - nhm$. Since $K_{\max} \leq h^m$, it is also a (K_{\max}, A) expander.

Note that the number of left-vertices in Γ is $q^n \geq N$, and the number of right-vertices is

$$M = q^{m+1} \leq q^2 \cdot h^{(1+\alpha) \cdot (m-1)} \leq q^2 \cdot K_{\max}^{1+\alpha}.$$

The degree is

$$D = q \leq h^{1+\alpha} = O(nk/\varepsilon)^{1+1/\alpha} = O((\log N)(\log K_{\max})/\varepsilon)^{1+1/\alpha}.$$

To see that the expansion factor $A = q - nhm \geq q - nhk$ is at least $(1 - \varepsilon)D = (1 - \varepsilon)q$, note that

$$nhk \leq (\varepsilon/2) \cdot h^{1+\alpha} \leq \varepsilon q,$$

where the first inequality holds because $h^\alpha \geq 2nk/\varepsilon$.

Finally, the construction is explicit because a description of \mathbb{F}_q for q a power of 2 (i.e. an irreducible polynomial of degree $\log q$ over \mathbb{F}_2) as well as an irreducible polynomial $E(Y)$ of degree n over \mathbb{F}_q can be found in time $\text{poly}(n, \log q) = \text{poly}(\log N, \log D)$. \square

These expanders are of polylogarithmic rather than constant degree. But the expansion is almost as large as possible given the degree ($A = (1 - \varepsilon) \cdot D$), and the size of the right-hand side is almost as small as possible (in a (K, A) expander, we must have $M \geq KA = (1 - \varepsilon)KD$). In particular, these expanders can be used in the data structure application of Problem 4.10 — storing a K -sized subset of $[N]$ using $K^{1.01} \cdot \text{polylog}(N)$ bits in such a way that membership can be probabilistically tested by reading 1 bit of the data structure. (An efficient solution to that application actually requires more than the graph being explicit in the usual sense, but also that there are efficient algorithms for finding all left-vertices having at least some δ fraction neighbors in a given set $T \subset [M]$ of right vertices, but the expanders above can be shown to have that property by a variant of the list-decoding algorithm above.)

A deficiency of the expander of Theorem 5.35 is that the size of the right-hand side is polynomial in K and D (for constant α), whereas the optimal bound is $M = O(KD/\varepsilon)$. Achieving the latter, while keeping the left-degree polylogarithmic, is an open problem:

Open Problem 5.36. Construct $(= K, A)$ bipartite expanders with N left-vertices, degree $D = \text{poly}(\log N)$, expansion $A = .99D$, and $M = O(KD)$ right-hand vertices.

We remark that a construction where D is *quasipolynomial* in $\log N$ is known.

5.5 Exercises

Problem 5.1 (Limits of List Decoding). Show that if there exists a q -ary code $\mathcal{C} \subset \Sigma^{\hat{n}}$ of rate ρ that is (δ, L) list-decodable, then $\rho \leq 1 - H_q(\delta, \hat{n}) + (\log_q L)/\hat{n}$

Problem 5.2 (Concatenated Codes). For codes $\text{Enc}_1 : \{1, \dots, N\} \rightarrow \Sigma_1^{n_1}$ and $\text{Enc}_2 : \Sigma_1 \rightarrow \Sigma_2^{n_2}$, their *concatenation* $\text{Enc} : \{1, \dots, N\} \rightarrow \Sigma_2^{n_1 n_2}$ is defined by

$$\text{Enc}(m) = \text{Enc}_2(\text{Enc}_1(m)_1)\text{Enc}_2(\text{Enc}_1(m)_2) \cdots \text{Enc}_2(\text{Enc}_1(m)_{n_1}).$$

This is typically used as a tool for reducing alphabet size, e.g. with $\Sigma_2 = \{0, 1\}$.

- (1) Prove that if Enc_1 has minimum distance δ_1 and Enc_2 has minimum distance δ_2 , then Enc has minimum distance at least $\delta_1 \delta_2$.

- (2) Prove that if Enc_1 is $(1 - \varepsilon_1, \ell_1)$ list-decodable and Enc_2 is (δ_2, ℓ_2) list-decodable, then Enc is $((1 - \varepsilon_1 \ell_2) \cdot \delta_2, \ell_1 \ell_2)$ list-decodable.
- (3) By concatenating a Reed–Solomon code and a Hadamard code, show that for every $n \in \mathbb{N}$ and $\varepsilon > 0$, there is a (fully) explicit code $\text{Enc} : \{0, 1\}^n \rightarrow \{0, 1\}^{\hat{n}}$ with blocklength $\hat{n} = O(n^2/\varepsilon^2)$ with minimum distance at least $1/2 - \varepsilon$. Furthermore, show that with blocklength $\hat{n} = \text{poly}(n, 1/\varepsilon)$, we can obtain a code that is $(1/2 - \varepsilon, \text{poly}(1/\varepsilon))$ list-decodable in *polynomial time*. (Hint: the inner code can be decoded by brute force.)
-

Problem 5.3 (List Decoding implies Unique Decoding for Random Errors). (1)

Suppose that $\mathcal{C} \subset \{0, 1\}^{\hat{n}}$ is a code with minimum distance at least $1/4$ and rate at most $\alpha\varepsilon^2$ for a constant $\alpha > 0$, and we transmit a codeword $c \in \mathcal{C}$ over a channel in which each bit is flipped with probability $1/2 - 2\varepsilon$. Show that if α is a sufficiently small constant (independent of \hat{n} and ε), then all but exponentially small probability over the errors, c will be the unique codeword at distance at most $1/2 - \varepsilon$ from the received word r .

- (2) Using Problem 5.2, deduce that for every $\varepsilon > 0$ and $n \in \mathbb{N}$, there is an explicit code of blocklength $\hat{n} = \text{poly}(n, 1/\varepsilon)$ that can be uniquely decoded from $(1/2 - 2\varepsilon)$ random errors as above in polynomial time.
-

Problem 5.4 (Linear Codes). For a prime power q , a q -ary code $\mathcal{C} \subset \mathbb{F}_q^{\hat{n}}$ is called *linear* if \mathcal{C} is a linear subspace of $\mathbb{F}_q^{\hat{n}}$. That is, for every $u, v \in \mathcal{C}$ and $\alpha \in \mathbb{F}_q$, we have $u + v \in \mathcal{C}$ and $\alpha u \in \mathcal{C}$.

- (1) Verify that the Hadamard, Reed–Solomon, and Reed–Muller codes are all linear.
- (2) Show that if \mathcal{C} is linear, then the minimum distance of \mathcal{C} equals the *minimum weight* of \mathcal{C} , which is defined to be $\min_{c \in \mathcal{C} \setminus \{0\}} |\{i : c_i \neq 0\}|/\hat{n}$.
- (3) Show that if \mathcal{C} is a subspace of dimension n , then its rate is n/\hat{n} and it has an encoding function $\text{Enc} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^{\hat{n}}$ that is a linear map. Moreover, the encoding function can be made to have the property that there is a set $S \subset [\hat{n}]$ of n coordinates such that $\text{Enc}(m)|_S = m$ for every $m \in \mathbb{F}_q^n$. (Here $c|_S$ is the projection of c onto the coordinates in S .) Such encodings are called *systematic*, and will be useful when we study locally decodable codes in Chapter 7.
- (4) Find explicit systematic encodings for the Hadamard and Reed–Solomon codes.
- (5) Show that the nonconstructive bound of Theorem 5.8, Part 1, can be achieved by a linear code. That is, for all prime powers q , integers $n, \hat{n} \in \mathbb{N}$, and $\delta \in (0, 1 - 1/q)$, there exists a *linear* q -ary code of block length \hat{n} , minimum distance at least δ , and rate at least $\rho = n/\hat{n}$ provided $\rho \leq 1 - H_q(\delta, \hat{n})$. (Hint: chose a random linear map $\text{Enc} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^{\hat{n}}$ and use Part 2.)
-

Problem 5.5 (LDPC Codes). Given a bipartite multigraph G with N left-vertices and M right-vertices, we can obtain a linear code $\mathcal{C} \subset \{0, 1\}^N$ (where we view $\{0, 1\}$ as the field of two elements) by:

$$\mathcal{C} = \{c \in \{0, 1\}^N : \forall j \in [M] \oplus_{i \in \Gamma(j)} c_i = 0\},$$

where $\Gamma(j)$ denotes the set of neighbors of vertex j . When G has small left-degree D (e.g. $D = O(1)$), then \mathcal{C} is called a *low-density parity check (LDPC) code*.

- (1) Show that \mathcal{C} has rate at least $1 - M/N$.
- (2) Show that if G is a (K, A) expander for $A > D/2$, then \mathcal{C} has minimum distance at least $\delta = K/N$.
- (3) Show that if G is a $(K, (1 - \varepsilon)D)$ expander for a sufficiently small constant ε , then \mathcal{C} has a polynomial-time δ -decoder for $\delta = (1 - 3\varepsilon) \cdot K/N$. Assume that G is given as input to the decoder. (Hint: given a received word $r \in \{0, 1\}^n$, flip all coordinates of r for which at least $2/3$ of the neighboring parity checks are not satisfied, and argue that the number of errors decreases by a constant factor. It may be useful to use the results of Problem 4.10.)

By a probabilistic argument like Theorem 4.4, graphs G as above exist with $D = O(1)$, $K = \Omega(N)$, $M = (1 - \Omega(1))N$, arbitrarily small constant $\varepsilon > 0$, and $N \rightarrow \infty$, and in fact explicit constructions are known. This yields explicit LDPC codes with constant rate and constant distance (“asymptotically good LDPC codes”).

- Problem 5.6 (Improved list decoding of Reed–Solomon Codes).** (1) Show that there is a polynomial-time algorithm for list decoding the Reed–Solomon codes of degree d over \mathbb{F}_q up to distance $1 - \sqrt{2d/q}$, improving the $1 - 2\sqrt{d/q}$ bound from Theorem 5.19. (Hint: do not use fixed upper bounds on the individual degrees of the interpolating polynomial $Q(Y, Z)$, but rather allow as many monomials as possible.)
- (2) (*) Improve the list-decoding radius further to $1 - \sqrt{d/q}$ by using the following “method of multiplicities”. First, require the interpolating polynomial $Q(Y, Z)$ to have a zero of multiplicity s at each point $(y, r(y))$ — that is, the polynomial $Q(Y + y, Z + r(y))$ should have no monomials of degree smaller than s . Second, use the fact that a univariate polynomial $R(Y)$ of degree t can have at most t roots, counting multiplicities.
-

Problem 5.7 (Twenty Questions). In the game of 20 questions, an oracle has an arbitrary secret $s \in \{0, 1\}^n$ and the aim is to determine the secret by asking the oracle as few yes/no questions about s as possible. It is easy to see that n questions are necessary and sufficient. Here we consider a variant where the oracle has two secrets $s_1, s_2 \in \{0, 1\}^n$, and can adversarially decide to answer each question according to either s_1 or s_2 . That is, for a question $f : \{0, 1\}^n \rightarrow \{0, 1\}$, the oracle may answer with either $f(s_1)$ or $f(s_2)$. Here it turns out to be impossible to pin down either of the secrets with certainty, no matter how many questions we ask, but we can hope to compute a small list L of secrets such that $|L \cap \{s_1, s_2\}| \neq 0$. (In fact, $|L|$ can be made as small as 2.) This variant of twenty questions apparently was motivated by problems in Internet traffic routing.

- (1) Let $\text{Enc} : \{0, 1\}^n \rightarrow \{0, 1\}^{\hat{n}}$ be a code such that every two codewords in Enc *agree* in at least a $1/2 - \varepsilon$ fraction of positions and that Enc has a polynomial-time $(1/4 + \varepsilon, \ell)$ list-decoding algorithm. Show how to solve the above problem in polynomial time by asking the \hat{n} questions $\{f_i\}$ defined by $f_i(x) = \text{Enc}(x)_i$.
- (2) Taking Enc to be the code constructed in Problem 1, deduce that $\hat{n} = \text{poly}(n)$ questions suffices.

5.6 Chapter Notes and References

Standard texts on coding theory include MacWilliams–Sloane [MS] and the Handbook of Coding Theory [PHB]. The lecture notes of Sudan [Sud2] present coding theory with an algorithmic perspective, and Guruswami [Gur2] gives a thorough treatment of list decoding.

The field of coding theory began with Shannon’s seminal 1948 paper [Sha3], which proposed the study of stochastic error models and proved that a random error-correcting code achieves an optimal rate-distance tradeoff with high probability. The study of decoding from worst-case errors began a couple of years later with the work of Hamming [Ham]. The notion of list decoding was proposed independently in the late 1950’s, in the work of Elias [Eli1] and Wozencraft [Woz].

The nonconstructive bound for the tradeoff between rate and minimum distance of Theorem 5.8 (Part 1) is due to Gilbert [Gil1], and its extension to linear codes in Problem 5.4 (Part 5) is due to Varshamov [Var]; together they are known as the Gilbert–Varshamov Bound. For more on the algebraic geometric codes and how they can be used to beat the Gilbert–Varshamov Bound, see the text by Stichtenoth [Sti1]. The nonconstructive bound for list decoding of Theorem 5.8 (Part 2) is due to Elias [Eli3], and it has been extended to linear codes in [ZP, GHSZ, GHK]. The Johnson Bound (Proposition 5.10) was proven for binary codes by Johnson [Joh2, Joh1]. An optimized form of the bound can be found in [GS3].

The binary ($q = 2$) case of Reed–Muller codes was introduced independently by Reed [Ree] and Muller [Mul] in the mid-50’s. Reed–Solomon codes were introduced in 1960 by Reed and Solomon [RS]. Polynomial time unique-decoding algorithms for Reed–Solomon Codes include those of Peterson [Pet] and Berlekamp [Ber2], . The first nontrivial list decoding algorithm was given by Goldreich and Levin [GL]; while their algorithm is stated in the language of “hardcore bits for one-way functions,” it can be viewed as an efficient “local” list-decoding algorithm for the Hadamard Code. (Local decoding algorithms are discussed in Chapter 7.) The list-decoding algorithm for Reed–Solomon Codes of Theorem 5.19 is from the seminal work of Sudan [Sud1], which sparked a resurgence in the study of list decoding. The improved decoding algorithm of Problem 5.6, Part 2 is due to Guruswami and Sudan [GS1]. Parvaresh–Vardy codes and their decoding algorithm are from [PV]. Folded Reed–Solomon Codes and their capacity-achieving list-decoding algorithm are due to Guruswami and Rudra [GR]. (In all cases, our presentation uses a simplified setting of parameters compared to the original algorithms.) The list size, decoding time, and alphabet size of Folded Reed–Solomon Codes have recently been improved in [Gur3, DL, GX], with the best parameters to date being achieved in the randomized construction of [GX], while [DL] give a fully deterministic construction.

The list-decoding views of expanders and samplers emerged out of work on randomness extractors (the subject of Chapter 6). Specifically, a close connection between extractors and expanders

was understood already at the time that extractors were introduced by Nisan and Zuckerman [NZ]. An equivalence between extractors and averaging samplers was established by Zuckerman [Zuc2] (building on previous connections between other types of samplers and expanders [Sip2, CW1]). A connection between list-decodable codes and extractors emerged in the work of Trevisan [Tre1], and the list-decoding view of extractors was crystallized by Ta-Shma and Zuckerman [TZ]. The list-decoding formulation of vertex expansion and the construction of expanders from Parvaresh–Vardy codes is due to Guruswami, Umans, and Vadhan [GUV].

Code concatenation (Problem 5.2) was introduced by Forney [For], who used it to construct explicit, efficiently decodable binary codes that achieve capacity for random errors. Efficient list-decoding algorithms for the concatenation of a Reed–Solomon code with a Hadamard code were given in [GS2]. Low-density parity check (LDPC) codes (Problem 5.5) were introduced by Gallager [Gal]. The use of expansion to analyze such codes and their decoding algorithm comes from Sipser and Spielman [SS]. Explicit expanders suitable for Problem 5.5 were given by [CRVW]. The surveys of Guruswami [Gur1, Gur2] describe the state of the art in expander-based constructions of codes and in LDPC codes, respectively.

The question of Problem 5.7 (Twenty Questions) was posed in [CGL], motivated by Internet traffic routing applications. The solution using list decoding is from [AGKS].

6

Randomness Extractors

Randomness extractors are functions that extract almost-uniform bits from sources of biased and correlated bits. The original motivation for extractors was to simulate randomized algorithms with weak random sources as might arise in nature. This motivation is still compelling, but extractors have taken on a much wider significance in the years since they were introduced. They have found numerous applications in theoretical computer science beyond this initial motivating one, in areas from cryptography to distributed algorithms to hardness of approximation. More importantly from the perspective of this survey, they have played a major unifying role in the theory of pseudorandomness. Indeed, the links between the various pseudorandom objects we are studying in this survey (expander graphs, randomness extractors, list-decodable codes, pseudorandom generators, samplers) were all discovered through work on extractors (even though now we find it more natural to present these links in the language of list decoding, as introduced in Section 5.3).

6.1 Motivation and Definition

6.1.1 Deterministic Extractors

Typically, when we design randomized algorithms or protocols, we assume that all algorithms/parties have access to sources of *perfect* randomness, i.e. bits that are unbiased and completely independent. However, when we implement these algorithms, the physical sources of randomness to which we have access may contain biases and correlations. For example, we may use low-order bits of the system clock, the user’s mouse movements, or a noisy diode based on quantum effects. While these sources may have some randomness in them, the assumption that the source is perfect is a strong one, and thus it is of interest to try and relax it.

Ideally, what we would like is a compiler that takes any algorithm A that works correctly when fed perfectly random bits U_m , and produces a new algorithm A' that will work even if it is fed random bits $X \in \{0, 1\}^n$ that come from a “weak” random source. For example, if A is a **BPP** algorithm, then we would like A' to also run in probabilistic polynomial time. One way to design such compilers is to design a *randomness extractor* $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ such that $\text{Ext}(X)$ is

distributed uniformly in $\{0, 1\}^m$.

IID-Bit Sources. A simple version of this question was already considered by von Neumann. He looked at sources that consist of boolean random variables $X_1, X_2, \dots, X_n \in \{0, 1\}$ that are independent but biased. That is, for every i , $\Pr[X_i = 1] = \delta$ for some unknown δ . How can such a source be converted into a source of independent, unbiased bits? Von Neumann proposed the following extractor: Break all the variables in pairs and for each pair output 0 if the outcome was 01, 1 if the outcome was 10, and skip the pair if the outcome was 00 or 11. This will yield an unbiased random bit after $1/(2\delta(1 - \delta))$ pairs on average.

Independent-Bit Sources. Lets now look at a bit more interesting class of sources in which all the variables are still independent but the bias is no longer the same. Specifically, for every i , $\Pr[X_i = 1] = \delta_i$ and $0 < \delta \leq \delta_i \leq 1 - \delta$ for some constant $\delta > 0$. How can we deal with such a source?

It can be shown that when we take a parity of ℓ bits from such an independent-bit source, the result approaches an unbiased coin flip exponentially fast in ℓ , i.e. $|\Pr[\bigoplus_{i=1}^{\ell} X_i = 1] - 1/2| = 2^{-\Omega(\ell)}$. The result is not a perfect coin flip but is as good as one for almost all purposes.

Let's be more precise about the problems we are studying. A *source* on $\{0, 1\}^n$ is simply a random variable X taking values in $\{0, 1\}^n$. In each of the above examples, there is an implicit *class* of sources being studied. For example, $\text{IndBits}_{n,\delta}$ is the class of sources X on $\{0, 1\}^n$ where the bits X_i are independent and satisfy $\delta \leq \Pr[X_i = 1] \leq 1 - \delta$. We could define $\text{IIDBits}_{n,\delta}$ to be the same with the further restriction that all of the X_i 's are identically distributed, i.e. $\Pr[X_i = 1] = \Pr[X_j = 1]$ for all i, j , thereby capturing von Neumann sources.

Definition 6.1 (deterministic extractors).¹ Let \mathcal{C} be a class of sources on $\{0, 1\}^n$. An ε -*extractor* for \mathcal{C} is a function $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ such that for every $X \in \mathcal{C}$, $\text{Ext}(X)$ is " ε -close" to U_m .

Note that we want a *single* function Ext that works for all sources in the class. This captures the idea that we do not want to assume we know the exact distribution of the physical source we are using, but only that it comes from some class. For example, for $\text{IndBits}_{n,\delta}$, we know that the bits are independent and none are too biased, but not the specific bias of each bit. Note also that we only allow the extractor *one* sample from the source X . If we want to allow multiple independent samples, then this should be modelled explicitly in our class of sources; ideally we would like to minimize the independence assumptions used.

We still need to define what we mean for the output to be ε -close to U_m .

Definition 6.2. For random variables X and Y taking values in \mathcal{U} , their *statistical difference* (also known as *variation distance*) is $\Delta(X, Y) = \max_{T \subset \mathcal{U}} |\Pr[X \in T] - \Pr[Y \in T]|$. We say that X and Y are ε -*close* if $\Delta(X, Y) \leq \varepsilon$.

¹Such extractors are called *deterministic* or *seedless* to contrast with the probabilistic or *seeded* randomness extractors we will see later.

Intuitively, any event in X happens in Y with the same probability $\pm\varepsilon$. This is perhaps the most natural measure of distance for probability distributions (much more so than the ℓ_2 distance we used in the study of random walks). In particular, it satisfies the following natural properties.

Lemma 6.3 (properties of statistical difference). Let $X, Y, Z, X_1, X_2, Y_1, Y_2$ be random variables taking values in a universe \mathcal{U} . Then,

- (1) $\Delta(X, Y) \geq 0$, with equality iff X and Y are identically distributed,
 - (2) $\Delta(X, Y) \leq 1$, with equality iff X and Y have disjoint supports,
 - (3) $\Delta(X, Y) = \Delta(Y, X)$,
 - (4) $\Delta(X, Z) \leq \Delta(X, Y) + \Delta(Y, Z)$,
 - (5) for every function f , we have $\Delta(f(X), f(Y)) \leq \Delta(X, Y)$,
 - (6) $\Delta((X_1, X_2), (Y_1, Y_2)) \leq \Delta(X_1, Y_1) + \Delta(X_2, Y_2)$ if X_1 and X_2 , as well as Y_1 and Y_2 , are independent, and
 - (7) $\Delta(X, Y) = \frac{1}{2} \cdot |X - Y|_1$, where $|\cdot|_1$ is the ℓ_1 distance. (Thus, X is ε -close to Y iff we can transform X into Y by “shifting” at most an ε fraction of probability mass.)
-

We now observe that extractors according to this definition give us the “compilers” we want.

Proposition 6.4. Let $A(w; r)$ be a randomized algorithm such that $A(w; U_m)$ has error probability at most γ , and let $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be an ε -extractor for a class \mathcal{C} of sources on $\{0, 1\}^n$. Define $A'(w; x) = A(w; \text{Ext}(x))$. Then for every source $X \in \mathcal{C}$, $A'(w; X)$ has error probability at most $\gamma + \varepsilon$.

This application identifies some additional properties we’d like from our extractors. We’d like the extractor itself to be efficiently computable (e.g. polynomial time). In particular, to get m almost-uniform bits out, we should need at most $n = \text{poly}(m)$ bits from the weak random source.

We can cast our earlier extractor for sources of independent bits in this language:

Proposition 6.5. For every constant $\delta > 0$, every $n, m \in \mathbb{N}$, there is a polynomial-time computable function $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ that is an ε -extractor for $\text{IndBits}_{n, \delta}$, with $\varepsilon = m \cdot 2^{-\Omega(n/m)}$.

In particular, taking $n = m^2$, we get exponentially small error with a source of polynomial length.

Proof. Ext breaks the source into m blocks of length $\lfloor n/m \rfloor$ and outputs the parity of each block. □

Unpredictable-Bit Sources (aka Santha–Vazirani Sources). Another interesting class of sources, which looks similar to the previous example is the class $\text{UnpredBits}_{n, \delta}$ of *unpredictable-bit sources*. These are the sources that for every i , every $x_1, \dots, x_n \in \{0, 1\}$ and some constant $\delta > 0$, satisfy

$$\delta \leq \Pr [X_i = 1 \mid X_1 = x_1, X_2 = x_2, \dots, X_{i-1} = x_{i-1}] \leq 1 - \delta$$

The parity extractor used above will be of no help with this source since the next bit could be chosen in a way that the parity will be equal to 1 with probability δ . Indeed, there does not exist any nontrivial extractor for these sources — the best we can do is output the first bit:

Proposition 6.6. For every $n \in \mathbb{N}$, $\delta > 0$, and fixed extraction function $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}$ there exists a source $X \in \text{UnpredBits}_{n, \delta}$ such that either $\Pr[\text{Ext}(X) = 1] \leq \delta$ or $\Pr[\text{Ext}(X) = 1] \geq 1 - \delta$. That is, there is no ε -extractor for $\text{UnpredBits}_{n, \delta}$ for $\varepsilon < 1/2 - \delta$.

The proof is left as an exercise (Problem 6.6).

Nevertheless, as we will see, the answer to the question whether we can simulate **BPP** algorithms with unpredictable sources will be “yes”! Indeed, we will even be able to handle a much more general class of sources, introduced in the next section.

6.1.2 Entropy Measures and General Weak Sources

Intuitively, to extract m almost-uniform bits from a source, the source must have at least “ m bits of randomness” in it. (In particular, its support cannot be much smaller than 2^m .) Ideally, this is all we would like to assume about a source. Thus, we need some measure of how much randomness is in a random variable; this can be done using various notions of *entropy* described below.

Definition 6.7 (entropy measures). Let X be a random variable. Then

- the *Shannon entropy* of X is:

$$H_{Sh}(X) = \mathbb{E}_{x \leftarrow X} \left[\log \frac{1}{\Pr[X = x]} \right].$$

- the *Rényi entropy* of X is:

$$H_2(X) = \log \left(\frac{1}{\mathbb{E}_{x \leftarrow X} [\Pr[X = x]]} \right) = \log \frac{1}{\text{CP}(X)}, \text{ and}$$

- the *min-entropy* of X is:

$$H_\infty(X) = \min_x \left\{ \log \frac{1}{\Pr[X = x]} \right\},$$

where all logs are base 2.

Rényi entropy $H_2(X)$ should not be confused with the binary entropy function $H_2(\delta)$ from Definition 5.6. Indeed, the q -ary entropy $H_q(\delta)$ is equal to the *Shannon* entropy of a random variable that equals 1 with probability $1 - \delta$ and is uniformly distributed in $\{2, \dots, q\}$ with probability δ .

All the three measures satisfy the following properties we would expect from a measure of randomness:

Lemma 6.8 (properties of entropy). For each of the entropy measures $H \in \{H_{Sh}, H_2, H_\infty\}$ and random variables X, Y , we have:

- $H(X) \geq 0$, with equality iff X is supported on a single element,
 - $H(X) \leq \log |\text{Supp}(X)|$, with equality iff X is uniform on $\text{Supp}(X)$,
 - if X, Y are independent, then $H((X, Y)) = H(X) + H(Y)$,
 - for every deterministic function f , we have $H(f(X)) \leq H(X)$, and
 - for every X , we have $H_\infty(X) \leq H_2(X) \leq H_{Sh}(X)$.
-

To illustrate the differences between the three notions, consider a source X such that $X = 0^n$ with probability 0.99 and $X = U_n$ with probability 0.01. Then $H_{Sh}(X) \geq 0.01n$ (contribution from the uniform distribution), $H_2(X) \leq \log(1/.99^2) < 1$ and $H_\infty(X) \leq \log(1/.99) < 1$ (contribution from 0^n). Note that even though X has Shannon entropy linear in n , we cannot expect to extract bits that are close to uniform or carry out any useful randomized computations with one sample from X , because it gives us nothing useful 99% of the time. Thus, we should use the stronger measures of entropy given by H_2 or H_∞ .

Then why is Shannon entropy so widely used in information theory results? The reason is that such results typically study what happens when you have many independent samples from the source (whereas we only allow one sample). In the case of many samples, it turns out that the source is “close” to one where the min-entropy is roughly equal to the Shannon entropy. Thus the distinction between these entropy measures becomes less significant. Moreover, Shannon entropy satisfies many nice identities that make it quite easy to work with. Min-entropy and Rényi entropy are much more delicate.

We will consider the task of extracting randomness from sources where all we know is a lower bound on the min-entropy:

Definition 6.9. A random variable X is a k -source if $H_\infty(X) \geq k$, i.e., if $\Pr[X = x] \leq 2^{-k}$.

A typical setting of parameters is $k = \delta n$ for some fixed δ , e.g., 0.01. We call δ the *min-entropy rate*. Some different ranges that are commonly studied (and are useful for different applications): $k = \text{polylog}(n)$, $k = n^\gamma$ for a constant $\gamma \in (0, 1)$, $k = \delta n$ for a constant $\delta \in (0, 1)$, and $k = n - O(1)$. The middle two ($k = n^\gamma$ and $k = \delta n$) are the most natural for simulating randomized algorithms with weak random sources.

Examples of k -sources:

- k random and independent bits, together with $n - k$ fixed bits (in an arbitrary order). These are called *oblivious bit-fixing sources*.
- k random and independent bits, and $n - k$ bits that depend arbitrarily on the first k bits. These are called *adaptive bit-fixing sources*.
- Unpredictable-bit sources with bias parameter δ . These are k -sources with $k = \log(1/(1 - \delta)^n) = \Theta(\delta n)$.
- Uniform distribution on a set $S \subset \{0, 1\}^n$ with $|S| = 2^k$. These are called *flat k -sources*.

It turns out that flat k -sources are really representative of general k -sources.

Lemma 6.10. Every k -source is a convex combination of flat k -sources (provided that $2^k \in \mathbb{N}$), i.e., $X = \sum p_i X_i$ with $0 \leq p_i \leq 1$, $\sum p_i = 1$ and all the X_i are flat k -sources.

That is, we can think of any k -source as being obtained by first selecting a flat k -source X_i according to some distribution (given by the p_i 's) and then selecting a random sample from X_i . This means that if we can compile probabilistic algorithms to work with flat k -sources, then we can compile them to work with any k -source.

Proof. Let X be a k -source on $[N]$. We can view X as partitioning a circle of unit circumference into N (half-open) intervals, where the t 'th interval has length exactly $\Pr[X = t]$. (If we associate the points on the circle with $[0, 1)$, then the t 'th interval is $[\Pr[X < t], \Pr[X \leq t])$.) Now consider a set S of K points spaced evenly on the circle. Then since each interval is half-open and has length at most $1/K$, each interval contains at most one point from S , so the uniform distribution on the set $T(S) = \{t : S \cap I_t \neq \emptyset\}$ is a flat k -source. Moreover, if we perform a uniformly random rotation of S on the circle to obtain a rotated set R and then choose a uniformly random element of $T(R)$, the probability that we output any value $t \in [N]$ is exactly the length of I_t , which equals $\Pr[X = t]$. Thus we have decomposed X as a convex combination of flat k -sources. (Specifically, $X = \sum_T p_T U_T$, where the sum is over subsets $T \subseteq [N]$ of size K , and $p_T = \Pr_R[T(R) = T]$.) \square

We also sketch another proof of Lemma 6.10 that can be more easily generalized to other classes of sources. We can view a random variable X taking values in $[N]$ as an N -dimensional vector, where $X(i)$ is the probability mass of i . Then X is a k -source if and only if $X(i) \in [0, 2^{-k}]$ for every $i \in [N]$ and $\sum_i X(i) = 1$. The set of vectors X satisfying these linear inequalities is a convex polytope in \mathbb{R}^N . By basic linear programming theory, all of the points in the polytope are convex combinations of its *vertices*, which are defined to be the points that make a maximal subset of the inequalities tight. By inspection, the vertices of the polytope of k -sources are those sources where $X(i) = 2^{-k}$ for 2^k values of i and $X(i) = 0$ for the remaining values of i ; these are simply the flat k -sources.

6.1.3 Seeded Extractors

Proposition 6.6 tells us that it is impossible to have deterministic extractors for unpredictable sources. Here we consider k -sources, which are more general than unpredictable sources, and hence it is also impossible to have deterministic extractors for them. The impossibility result for k -sources is stronger and simpler to prove.

Proposition 6.11. For any $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}$ there exists an $(n-1)$ -source X so that $\text{Ext}(X)$ is constant.

Proof. There exists $b \in \{0, 1\}$ so that $|\text{Ext}^{-1}(b)| \geq 2^n/2 = 2^{n-1}$. Then let X be the uniform distribution on $\text{Ext}^{-1}(b)$. \square

On the other hand, if we reverse the order of quantifiers, allowing the extractor to depend on the source, it is easy to see that good extractors exist and in fact a randomly chosen function will be a good extractor with high probability.

Proposition 6.12. For every $n, k, m \in \mathbb{N}$, every $\varepsilon > 0$, and every flat k -source X , if we choose a random function $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ with $m = k - 2 \log(1/\varepsilon) - O(1)$, then $\text{Ext}(X)$ will be ε -close to U_m with probability $1 - 2^{-\Omega(K\varepsilon^2)}$, where $K = 2^k$.

(In this chapter, we will make extensive use of the convention that capital variables are 2 raised to the power of the corresponding lowercase variable, such as $K = 2^k$ above.)

Proof. Choose our extractor randomly. We want it to have following property: for all $T \subset [M]$, $|\Pr[\text{Ext}(X) \in T] - \Pr[U_m \in T]| \leq \varepsilon$. Equivalently, $|\{x \in \text{Supp}(X) : \text{Ext}(x) \in T\}|/K$ differs from the density $\mu(T)$ by at most ε . For each point $x \in \text{Supp}(X)$, the probability that $\text{Ext}(x) \in T$ is $\mu(T)$, and these events are independent. By the Chernoff Bound (Theorem 2.21) for each fixed T , this condition holds with probability at least $1 - 2^{-\Omega(K\varepsilon^2)}$. Then the probability that condition is violated for at least one T is at most $2^M 2^{-\Omega(K\varepsilon^2)}$, which is less than 1 for $m = k - 2 \log(1/\varepsilon) - O(1)$. \square

Note that the failure probability is doubly-exponentially small in k . Naively, one might hope that we could get an extractor that's good for all flat k -sources by a union bound. But the number of flat k -sources is $\binom{N}{K} \approx N^K$ (where $N = 2^n$), which is unfortunately a larger double-exponential in k . We can overcome this gap by allowing the extractor to be “slightly” probabilistic, i.e. allowing the extractor a *seed* consisting of a small number of truly random bits in addition to the weak random source. We can think of this seed of truly random bits as a random choice of an extractor from family of extractors. This leads to the following crucial definition:

Definition 6.13 (seeded extractors). A function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε) -*extractor* if for every k -source X on $\{0, 1\}^n$, $\text{Ext}(X, U_d)$ is ε -close to U_m .

(Sometimes we will refer to extractors $\text{Ext} : [N] \times [D] \rightarrow [M]$ whose domain and range do not consist of bit-strings. These are defined in the natural way, requiring that $\text{Ext}(X, U_{[D]})$ is ε -close to $U_{[M]}$.)

The goal is to construct extractors that minimize d and maximize m . We prove the following theorem.

Theorem 6.14. For every $n \in \mathbb{N}$, $k \in [0, n]$ and $\varepsilon > 0$, there exists a (k, ε) -extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $m = k + d - 2 \log(1/\varepsilon) - O(1)$ and $d = \log(n - k) + 2 \log(1/\varepsilon) + O(1)$.

One setting of parameters to keep in mind (for our application of simulating randomized algorithms with a weak source) is $k = \delta n$, with δ a fixed constant (e.g. $\delta = 0.01$), and ε a fixed constant (e.g. $\varepsilon = 0.01$).

Proof. We use the Probabilistic Method. By Lemma 6.10, it suffices for Ext to work for flat k -sources. Choose the extractor Ext at random. Then the probability that the extractor fails is at most the number of flat k -sources times the probability Ext fails for a fixed flat k -source. By the above proposition, the probability of failure for a fixed flat k -source is at most $2^{-\Omega(KD\varepsilon^2)}$, since

(X, U_d) is a flat $(k + d)$ -source) and $m = k + d - 2 \log(\frac{1}{\varepsilon}) - O(1)$. Thus the total failure probability is at most

$$\binom{N}{K} \cdot 2^{-\Omega(KD\varepsilon^2)} \leq \left(\frac{Ne}{K}\right)^K 2^{-\Omega(KD\varepsilon^2)}.$$

The latter expression is less than 1 if $D\varepsilon^2 \geq c \log(Ne/K) = c \cdot (n - k) + c'$ for constants c, c' . This is equivalent to $d \geq \log(n - k) + 2 \log(\frac{1}{\varepsilon}) + O(1)$. \square

It turns out that both bounds (on m and d) are individually tight up to the $O(1)$ terms.

Recall that our motivation for extractors was to simulate randomized algorithms given *only* a weak random source, so allowing a truly random seed may seem to defeat the purpose. However, if the seed is of logarithmic length as in Theorem 6.14, then instead of selecting it randomly, we can enumerate all possibilities for the seed and take a majority vote.

Proposition 6.15. Let $A(w; r)$ be a randomized algorithm for computing a function f such that $A(w; U_m)$ has error probability at most γ (i.e. $\Pr[A(w; U_m) \neq f(w)] \leq \gamma$), and let $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ be a (k, ε) -extractor. Define

$$A'(w; x) = \text{maj}_{y \in \{0, 1\}^d} \{A(w; \text{Ext}(x, y))\}.$$

Then for every k -source X on $\{0, 1\}^n$, $A'(w; X)$ has error probability at most $2 \cdot (\gamma + \varepsilon)$.

Proof. The probability that $A(w; \text{Ext}(X, U_d))$ is incorrect is not more than the probability $A(w; U_m)$ is incorrect plus ε , i.e. $\gamma + \varepsilon$, by the definition of statistical difference. Then the probability that $\text{maj}_y A(w, \text{Ext}(X, y))$ is incorrect is at most $2 \cdot (\gamma + \varepsilon)$, because each error of $\text{maj}_y A(w; \text{Ext}(x, y))$ corresponds to $A(w; \text{Ext}(x, U_d))$ erring with probability at least $1/2$. \square

Note that the enumeration incurs a 2^d factor slowdown in the simulation. Thus, to retain running time $\text{poly}(m)$, we want to construct extractors where (a) $d = O(\log n)$; (b) Ext is computable in polynomial time; and (c) $m = n^{\Omega(1)}$.

We remark that the error probability in Proposition 6.15 can actually be made exponentially small by using an extractor that is designed for slightly lower min-entropy. (See Problem 6.2.)

We note that even though seeded extractors suffice for simulating randomized algorithms with only a weak source, they do not suffice for all applications of randomness in theoretical computer science. The trick of eliminating the random seed by enumeration does not work, for example, in cryptographic applications of randomness. Thus the study of deterministic extractors for restricted classes of sources remains a very interesting and active research direction. We, however, will focus on seeded extractors, due to their many applications and their connections to the other pseudorandom objects we are studying.

6.2 Connections to Other Pseudorandom Objects

As mentioned earlier, extractors have played a unifying role in the theory of pseudorandomness, through their close connections with a variety of other pseudorandom objects. In this section, we will see two of these connections. Specifically, how by reinterpreting them appropriately, extractors can be viewed as providing families of hash functions, and as being a certain type of highly expanding graphs.

6.2.1 Extractors as Hash Functions

Proposition 6.12 says that for any subset $S \subset [N]$ of size K , if we choose a completely random hash function $h : [N] \rightarrow [M]$ for $M \ll K$, then h will map the elements of S almost-uniformly to $[M]$. Equivalently, if we let H be distributed uniformly over all functions $h : [N] \rightarrow [M]$ and X be uniform on the set S , then $(H, H(X))$ is statistically close to $(H, U_{[M]})$, where we use the notation U_T to denote the uniform distribution on a set T . Can we use a smaller family of hash functions than the set of all functions $h : [N] \rightarrow [M]$? This gives rise to the following variant of extractors.

Definition 6.16 (strong extractors). Extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a *strong* (k, ε) -extractor if for every k -source X on $\{0, 1\}^n$, $(U_d, \text{Ext}(X, U_d))$ is ε -close to (U_d, U_m) . Equivalently, $\text{Ext}'(x, y) = (y, \text{Ext}(x, y))$ is a standard (k, ε) -extractor.

The nonconstructive existence proof of Theorem 6.14 can be extended to establish the existence of very good strong extractors:

Theorem 6.17. For every $n, k \in \mathbb{N}$ and $\varepsilon > 0$ there exists a strong (k, ε) -extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $m = k - 2 \log(1/\varepsilon) - O(1)$ and $d = \log(n - k) + 2 \log(1/\varepsilon) + O(1)$.

Note that the output length is $m \approx k$ instead of $m \approx k + d$; intuitively a strong extractor needs to extract randomness that is *independent* of the seed and thus can only get the k bits from the source.

We see that strong extractors can be viewed as very small families of hash functions having the almost-uniform mapping property mentioned above. Indeed, our first explicit construction of extractors is obtained by using pairwise independent hash functions.

Theorem 6.18 (Leftover Hash Lemma). If $\mathcal{H} = \{h : \{0, 1\}^n \rightarrow \{0, 1\}^m\}$ is a pairwise independent (or even 2-universal) family of hash functions where $m = k - 2 \log(1/\varepsilon)$, then $\text{Ext}(x, h) \stackrel{\text{def}}{=} h(x)$ is a strong (k, ε) -extractor. Equivalently, $\text{Ext}(x, h) = (h, h(x))$ is a standard (k, ε) -extractor.

Note that the seed length equals the number of random bits required to choose $h \stackrel{\text{R}}{\leftarrow} \mathcal{H}$, which is at least n by Problem 3.5.² This is far from optimal; for the purposes of simulating randomized algorithms we would like $d = O(\log n)$. However, the output length of the extractor is $m = k - 2 \log(1/\varepsilon)$, which is optimal up to an additive constant.

Proof. Let X be an arbitrary k -source on $\{0, 1\}^n$, \mathcal{H} as above, and $H \stackrel{\text{R}}{\leftarrow} \mathcal{H}$. Let d be the seed length. We show that $(H, H(X))$ is ε -close to $U_d \times U_m$ in the following three steps:

- (1) We show that the collision probability of $(H, H(X))$ is close to that of $U_d \times U_m$.
- (2) We note that this is equivalent to saying that the ℓ_2 distance between $(H, H(X))$ and $U_d \times U_m$ is small.

²Problem 3.5 refers to pairwise independent families, but a similar argument shows that universal families require $\Omega(n)$ random bits. (Instead of constructing orthogonal vectors, we construct vectors that have nonpositive dot product.)

(3) Then we deduce that the statistical difference is small, by recalling that the statistical difference equals half of the ℓ_1 distance, which can be (loosely) bounded by the ℓ_2 distance.

Proof of (1): By definition, $\text{CP}(H, H(X)) = \Pr[(H, H(X)) = (H', H'(X'))]$, where (H', X') is independent of and identically distributed to (H, X) . Note that $(H, H(X)) = (H', H'(X))$ if and only if $H = H'$ and either $X = X'$ or $X \neq X'$ but $H(X) = H(X')$. Thus

$$\begin{aligned} \text{CP}(H, H(X)) &= \text{CP}(H) \cdot \left(\text{CP}(X) + \Pr[H(X) = H(X') \mid X \neq X'] \right) \\ &\leq \frac{1}{D} \cdot \left(\frac{1}{K} + \frac{1}{M} \right) \leq \frac{1 + \varepsilon^2}{DM}. \end{aligned}$$

To see the penultimate inequality, note that $\text{CP}(H) = 1/D$ because there are D hash functions, $\text{CP}(X) \leq 1/K$ because $H_\infty(X) \geq k$, and $\Pr[H(X) = H(X') \mid X \neq X'] \leq 1/M$ by 2-universality.

Proof of (2):

$$\begin{aligned} \|(H, H(X)) - U_d \times U_m\|^2 &= \text{CP}(H, H(X)) - \frac{1}{DM} \\ &\leq \frac{1 + \varepsilon^2}{DM} - \frac{1}{DM} = \frac{\varepsilon^2}{DM}. \end{aligned}$$

Proof of (3): Recalling that the statistical difference between two random variables X and Y is equal to $\frac{1}{2}|X - Y|_1$, we have:

$$\begin{aligned} \Delta((H, H(X)), U_d \times U_m) &= \frac{1}{2} \cdot |(H, H(X)) - U_d \times U_m|_1 \\ &\leq \frac{\sqrt{DM}}{2} \cdot \|(H, H(X)) - U_d \times U_m\| \\ &\leq \frac{\sqrt{DM}}{2} \cdot \sqrt{\frac{\varepsilon^2}{DM}} \\ &= \frac{\varepsilon}{2}. \end{aligned}$$

Thus, we have in fact obtained a strong $(k, \varepsilon/2)$ -extractor. \square

The proof above actually shows that $\text{Ext}(x, h) = h(x)$ extracts with respect to collision probability, or equivalently, with respect to the ℓ_2 -norm. This property may be expressed in terms of Rényi entropy $H_2(Z) \stackrel{\text{def}}{=} \log(1/\text{CP}(Z))$. Indeed, we can define $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ to be a (k, ε) Rényi-entropy extractor if $H_2(X) \geq k$ implies $H_2(\text{Ext}(X, U_d)) \geq m - \varepsilon$ (or $H_2(U_d, \text{Ext}(X, U_d)) \geq m + d - \varepsilon$ for strong Rényi-entropy extractors). Then the above proof shows that pairwise-independent hash functions yield strong Rényi-entropy extractors.

In general, it turns out that an extractor with respect to Rényi entropy must have seed length $d \geq \min\{m/2, n - k\} - O(1)$ (as opposed to $d = O(\log n)$); this explains why the seed length in the above extractor is large. (See Problem 6.4.)

6.2.2 Extractors vs. Expanders

Extractors have a natural interpretation as graphs. Specifically, we can interpret an extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ as the neighbor function of a bipartite multigraph $G = ([N], [M], E)$

with $N = 2^n$ left-vertices, $M = 2^m$ right-vertices, and left-degree $D = 2^d$,³ where the r 'th neighbor of left-vertex u is $\text{Ext}(u, r)$. Typically $n \gg m$, so the graph is very unbalanced. It turns out that the extraction property of Ext is related to various “expansion” properties of G . In this section, we explore this relationship.

Let $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ be a (k, ε) -extractor and $G = ([N], [M], E)$ the associated graph. Recall that it suffices to examine Ext with respect to *flat* k -sources: in this case, the extractor property says that given a subset S of size $K = 2^k$ on the left, a random neighbor of a random element of S should be close to uniform on the right. In particular, if $S \subset [N]$ is a subset on the left of size K , then $|N(S)| \geq (1 - \varepsilon)M$. This property is just like vertex expansion, except that it ensures expansion only for sets of size exactly K , not any size $\leq K$. Recall that we call such a graph an $(= K, A)$ *vertex expander* (Definition 5.32). Indeed, this gives rise to the following weaker variant of extractors.

Definition 6.19 (dispersers). A function $\text{Disp} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε) -*disperser* if for every k -source X on $\{0, 1\}^n$, $\text{Disp}(X, U_d)$ has a support of size at least $(1 - \varepsilon) \cdot 2^m$.

While extractors can be used to simulate **BPP** algorithms with a weak random source, dispersers can be used to simulate **RP** algorithms with a weak random source. (See Problem 6.2.)

Then, we have:

Proposition 6.20. Let $n, m, d \in \mathbb{N}$, $K = 2^k \in \mathbb{N}$, and $\varepsilon > 0$. A function $\text{Disp} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε) -disperser iff the corresponding bipartite multigraph $G = ([N], [M], E)$ with left-degree D is an $(= K, A)$ vertex expander for $A = (1 - \varepsilon) \cdot M/K$.

Note that extractors and dispersers are interesting even when $M \ll K$, so the expansion parameter A may be less than 1. Indeed, $A < 1$ is interesting for vertex “expanders” when the graph is highly imbalanced. Still, for an *optimal* extractor, we have $M = \Theta(\varepsilon^2 KD)$ (because $m = k + d - 2 \log(1/\varepsilon) - \Theta(1)$), which corresponds to expansion factor $A = \Theta(\varepsilon^2 D)$. (An optimal disperser actually gives $A = \Theta(D/\log(1/\varepsilon))$.) Note this is smaller than the expansion factor of $D/2$ in Ramanujan graphs and $D - O(1)$ in random graphs; the reason is that those expansion factors are for “small” sets, whereas here we are asking for sets to expand to almost the entire right-hand side.

Now let’s look for a graph-theoretic property that is *equivalent* to the extraction property. Ext is a (k, ε) -extractor iff for every set $S \subset [N]$ of size K ,

$$\Delta(\text{Ext}(U_S, U_{[D]}), U_{[M]}) = \max_{T \subset [M]} \left| \Pr [\text{Ext}(U_S, U_{[D]}) \in T] - \Pr [U_{[M]} \in T] \right| \leq \varepsilon,$$

where U_S denotes the uniform distribution on S . This inequality may be expressed in graph-theoretic

³This connection is the reason we use d to denote the seed length of an extractor.

terms as follows. For every set $T \subset [M]$,

$$\begin{aligned} & \left| \Pr [\text{Ext}(U_S, U_{[D]}) \in T] - \Pr [U_{[M]} \in T] \right| \leq \varepsilon \\ & \Leftrightarrow \left| \frac{e(S, T)}{|S|D} - \frac{|T|}{M} \right| \leq \varepsilon \\ & \Leftrightarrow \left| \frac{e(S, T)}{ND} - \mu(S)\mu(T) \right| \leq \varepsilon\mu(S), \end{aligned}$$

where $e(S, T)$ is the number of edges from S to T (as in Definition 4.13).

Thus, we have:

Proposition 6.21. A function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε) -extractor iff the corresponding bipartite multigraph $G = ([N], [M], E)$ with left-degree D has the property that $|e(S, T)/ND - \mu(S)\mu(T)| \leq \varepsilon\mu(S)$ for every $S \subset [N]$ of size K and every $T \subset [M]$.

Note that this is very similar to the Expander Mixing Lemma (Lemma 4.15), which states that if a graph G has spectral expansion λ , then for *all* sets $S, T \subset [N]$ we have

$$\left| \frac{e(S, T)}{ND} - \mu(S)\mu(T) \right| \leq \lambda\sqrt{\mu(S)\mu(T)}.$$

It follows that if $\lambda\sqrt{\mu(S)\mu(T)} \leq \varepsilon\mu(S)$ for all $S \subset [N]$ of size K and all $T \subset [N]$, then G gives rise to a (k, ε) -extractor (by turning G into a D -regular bipartite graph with N vertices on each side in the natural way). It suffices for $\lambda \leq \varepsilon \cdot \sqrt{K/N}$ for this to work.

We can use this connection to turn our explicit construction of spectral expanders into an explicit construction of extractors. To achieve $\lambda \leq \varepsilon \cdot \sqrt{K/N}$, we can take an appropriate power of a constant-degree expander. Specifically, if G_0 is a D_0 -regular expander on N vertices with bounded second eigenvalue, we can consider the t th power of G_0 , $G = G_0^t$, where $t = O(\log((1/\varepsilon)\sqrt{N/K})) = O(n - k + \log(1/\varepsilon))$. The degree of G is $D = D_0^t$, so $d = \log D = O(t)$. This yields the following result:

Theorem 6.22. For every $n, k \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit (k, ε) -extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^n$ with $d = O(n - k + \log(1/\varepsilon))$.

Note that the seed length is significantly better than in the construction from pairwise-independent hashing when k is close to n , say $k = n - o(n)$ (i.e. $K = N^{1-o(1)}$). The output length is n , which is much larger than the typical output length for extractors (usually $m \ll n$). Using a Ramanujan graph (rather than an arbitrary constant-degree expander), the seed length can be improved to $d = n - k + 2\log(1/\varepsilon) + O(1)$, which yields an optimal output length $n = k + d - 2\log(1/\varepsilon) - O(1)$.

Another way of proving Theorem 6.22 is to use the fact that a random step on an expander decreases the ℓ_2 distance to uniform, like in the proof of the Leftover Hash Lemma. This analysis shows that we actually get a Rényi-entropy extractor; and thus explains the large seed length $d \approx n - k$.

Table 6.2.2 summarizes the main differences between “classic” expanders and extractors.

Expanders	Extractors
Measured by vertex or spectral expansion	Measured by min-entropy/statistical difference
Typically constant degree	Typically logarithmic or poly-logarithmic degree
All sets of size <i>at most</i> K expand	All sets of size <i>exactly</i> (or at least) K expand
Typically balanced	Typically unbalanced, bipartite graphs

Table 6.1 Differences between “classic” expanders and extractors

6.2.3 List Decoding View of Extractors

In this section, we cast extractors into the same list-decoding framework that we used to capture list-decodable codes, samplers, and expanders (in Section 5.3). Recall that all of these objects could be syntactically described as functions $\Gamma : [N] \times [D] \rightarrow [M]$, and their properties could be captured by bounding the sizes of sets of the form $\text{LIST}_\Gamma(T, \varepsilon) \stackrel{\text{def}}{=} \{x : \Pr_y[\Gamma(x, y) \in T] > \varepsilon\}$ for $T \subset [M]$. We also considered a generalization to functions $f : [M] \rightarrow [0, 1]$ where we defined $\text{LIST}_\Gamma(f, \varepsilon) \stackrel{\text{def}}{=} \{x : \mathbb{E}_y[f(\Gamma(x, y))] > \varepsilon\}$.

Conveniently, an extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ already meets the syntax of a function $\Gamma : [N] \times [D] \rightarrow [M]$ (matching our convention that $N = 2^n$, $D = 2^d$, $M = 2^m$). The extraction property can be described in the list-decoding framework as follows:

Proposition 6.23. Let $\Gamma = \text{Ext} : [N] \times [D] \rightarrow [M]$, let $K = 2^k \in \mathbb{N}$, and $\varepsilon \in [0, 1]$.

- (1) If Ext is a (k, ε) extractor, then for every $f : [M] \rightarrow [0, 1]$, we have

$$|\text{LIST}_\Gamma(f, \mu(f) + \varepsilon)| < K.$$

- (2) Suppose that for every $T \subset [M]$, we have

$$|\text{LIST}_\Gamma(T, \mu(T) + \varepsilon)| \leq K.$$

Then Ext is a $(k + \log(1/\varepsilon), 2\varepsilon)$ extractor.

Proof. (1) Suppose for contradiction that $|\text{LIST}_\Gamma(f, \mu(f) + \varepsilon)| \geq K$. Let X be uniformly distributed over $\text{LIST}_\Gamma(f, \mu(f) + \varepsilon)$. Then X is a k -source, and

$$\begin{aligned} \mathbb{E}[f(\text{Ext}(X, U_{[D]}))] &= \mathbb{E}_{x \stackrel{\text{R}}{\leftarrow} X} [f(\text{Ext}(x, U_{[D]}))] \\ &> \mu(f) + \varepsilon \\ &= \mathbb{E}[f(U_{[M]})] + \varepsilon. \end{aligned}$$

By Problem 6.1, this implies that $\text{Ext}(X, U_{[D]})$ and $U_{[M]}$ are ε -far, contradicting the hypothesis that Ext is a (k, ε) extractor.

- (2) Let X be any $(k + \log(1/\varepsilon))$ -source. We need to show that $\text{Ext}(X, U_{[D]})$ is 2ε -close to $U_{[M]}$. That is, we need to show that for every $T \subset [M]$, $\Pr[\text{Ext}(X, U_{[D]}) \in T] \leq \mu(T) + 2\varepsilon$.

So let T be any subset of $[M]$. Then

$$\begin{aligned}
& \Pr[\text{Ext}(X, U_{[D]}) \in T] \\
& \leq \Pr[X \in \text{LIST}(T, \mu(T) + \varepsilon)] + \Pr[\text{Ext}(X, U_{[D]}) \in T | X \notin \text{LIST}(T, \mu(T) + \varepsilon)] \\
& \leq |\text{LIST}(T, \mu(T) + \varepsilon)| \cdot 2^{-(k+\log(1/\varepsilon))} + (\mu(T) + \varepsilon) \\
& \leq K \cdot 2^{-(k+\log(1/\varepsilon))} + \mu(T) + \varepsilon \\
& = \mu(T) + 2\varepsilon.
\end{aligned}$$

□

The proposition does not give an *exact* list-decoding characterization of extractors, as the two parts are not exactly converses of each other. One difference is the extra $\log(1/\varepsilon)$ bits of entropy and the factor of 2 in ε appearing in Part 2. These are typically insignificant differences for extractors; indeed even an optimal extractor loses $\Theta(\log(1/\varepsilon))$ bits of entropy (cf. Theorem 6.14). A second difference is that Part 2 shows that extractors imply bounds on $|\text{LIST}(f, \mu(f) + \varepsilon)|$ even for fractional functions f , whereas Part 1 only requires bounds on $|\text{LIST}(T, \mu(T) + \varepsilon)|$ to deduce the extractor property. This only makes the result stronger, and indeed we will utilize this below.

Notice that the conditions characterizing extractors here are *identical* to the ones characterizing averaging samplers in Proposition 5.30. Actually, the condition in Part 1 is the one characterizing averaging samplers, whereas the condition in Part 2 is the one characterizing boolean averaging samplers. Thus we have:

Corollary 6.24. Let $\text{Ext} : [N] \times [D] \rightarrow [M]$ and $\text{Samp} : [N] \rightarrow [M]^D$ be such that $\text{Ext}(x, y) = \text{Samp}(x)_y$. Then:

- (1) If Ext is a (k, ε) extractor, then Samp is a $(K/N, \varepsilon)$ averaging sampler, where $K = 2^k$.
 - (2) If Samp is a $(K/N, \varepsilon)$ boolean averaging sampler, then Ext is a $(k + \log(1/\varepsilon), 2\varepsilon)$ extractor.
 - (3) If Samp is a (δ, ε) boolean averaging sampler, then Samp is a $(\delta/\varepsilon, 2\varepsilon)$ averaging sampler.
-

Thus, the only real difference between extractors and averaging samplers is one of perspective, and both perspectives can be useful. For example, in samplers, we measure the error probability $\delta = K/N = 2^k/2^n$, whereas in extractors we measure the min-entropy threshold k on its own. Thus, the sampler perspective can be more natural when δ is relatively large compared to $1/N$, and the extractor perspective when δ becomes quite close to $1/N$. Indeed, an extractor for min-entropy $k = o(n)$ corresponds to a sampler with error probability $\delta = 1/2^{(1-o(1))n}$, which means that each of the n bits of randomness used by the sampler reduces the error probability by almost a factor of 2!

We can now also describe the close connection between strong extractors and list-decodable codes when the alphabet size/output length is small.

Proposition 6.25. Let $\text{Ext} : [N] \times [D] \rightarrow [M]$ and $\text{Enc} : [N] \rightarrow [M]^D$ be such that $\text{Ext}(x, y) = \text{Enc}(x)_y$, and let $K = 2^k \in \mathbb{N}$.

- (1) If Ext is a strong (k, ε) extractor, then Enc is $(1 - 1/M - \varepsilon, K)$ list-decodable.

- (2) If Enc is $(1 - 1/M - \varepsilon, K)$ list-decodable, then Ext is a $(k + \log(1/\varepsilon), M \cdot \varepsilon)$ strong extractor.

Proof. (1) Follows from Corollaries 5.31 and 6.24.

- (2) Let X be a $(k + \log(1/\varepsilon))$ -source and $Y = U_{[D]}$. Then the statistical difference between $(Y, \text{Ext}(X, Y))$ and $Y \times U_{[M]}$ equals

$$\begin{aligned} \Delta((Y, \text{Ext}(X, Y)), Y \times U_{[M]}) &= \mathbb{E}_{y \leftarrow Y} [\Delta(\text{Ext}(X, y), U_{[M]})] \\ &= \mathbb{E}_{y \leftarrow Y} [\Delta(\text{Enc}(X)_y, U_{[M]})] \\ &\leq \frac{M}{2} \mathbb{E}_{y \leftarrow Y} \left[\max_z \Pr[\text{Enc}(X)_y = z] - 1/M \right] \end{aligned}$$

where the last inequality follows from the ℓ_1 formulation of statistical difference. So if we define $r \in [M]^D$ by setting r_y to be the value z maximizing $\Pr[\text{Enc}(X)_y = z] - 1/M$, we have:

$$\begin{aligned} \Delta((Y, \text{Ext}(X, Y)), Y \times U_{[M]}) &\leq \frac{M}{2} \cdot (\Pr[(Y, \text{Enc}(X)_Y) \in T_r] - 1/M), \\ &\leq \frac{M}{2} \cdot (\Pr[X \in \text{LIST}(T_r, 1/M + \varepsilon)] + \varepsilon) \\ &\leq \frac{M}{2} \cdot (2^{-(k + \log(1/\varepsilon))} \cdot K + \varepsilon) \\ &\leq M \cdot \varepsilon. \end{aligned}$$

□

Note that the quantitative relationship between extractors and list-decodable codes given by Proposition 6.25 deteriorates extremely fast as the output length/alphabet size increases. Nevertheless, the list-decoding view of extractors as given in Proposition 6.23 turns out to be quite useful.

6.3 Constructing Extractors

In the previous sections, we have seen that very good extractors exist — extracting almost all of the min-entropy from a source with only a logarithmic seed length. But the explicit constructions we have seen (via universal hashing and spectral expanders) are still quite far from optimal in seed length, and in particular cannot be used to give a polynomial-time simulation of **BPP** with a weak random source.

Fortunately, much better extractor constructions are known — ones that extract any constant fraction of the min-entropy using a logarithmic seed length, or extract all of the min-entropy using a polylogarithmic seed length. In this section, we will see how to construct such extractors.

6.3.1 Block Sources

We introduce a useful model of sources that has more structure than an arbitrary k -source:

Definition 6.26. A random variable $X = (X_1, X_2, \dots, X_t)$ is a (k_1, k_2, \dots, k_t) *block source* if for every x_1, \dots, x_{i-1} , $X_i |_{X_1=x_1, \dots, X_{i-1}=x_{i-1}}$ is a k_i -source. If $k_1 = k_2 = \dots = k_t = k$, then we call X a $t \times k$ *block source*.

Note that a (k_1, k_2, \dots, k_t) block source is also a $(k_1 + \dots + k_t)$ -source, but it comes with additional structure — each block is guaranteed to contribute some min-entropy. Thus, extracting randomness from block sources is an easier task than extracting from general sources.

The study of block sources has a couple of motivations.

- They are a natural and plausible model of sources in their own right. Indeed, they are more general than unpredictable-bit sources of Section 6.1.1: if $X \in \text{UnpredBits}_{n,\delta}$ is broken into t blocks of length $\ell = n/t$, then the result is a $t \times \delta'\ell$ block source, where $\delta' = \log(1/(1 - \delta))$.
- We can construct extractors for general weak sources by converting a general weak source into a block source. We will see how to do this later in the section.

We now illustrate how extracting from block sources is easier than from general sources. The idea is that we can extract almost-uniform bits from later blocks that are essentially independent of earlier blocks, and hence use these as a seed to extract more bits from the earlier blocks. Specifically, for the case of two blocks we have the following:

Lemma 6.27. Let $\text{Ext}_1 : \{0, 1\}^{n_1} \times \{0, 1\}^{d_1} \rightarrow \{0, 1\}^{m_1}$ be a (k_1, ε_1) -extractor, and $\text{Ext}_2 : \{0, 1\}^{n_2} \times \{0, 1\}^{d_2} \rightarrow \{0, 1\}^{m_2}$ be a (k_2, ε_2) -extractor with $m_2 \geq d_1$. Define $\text{Ext}'((x_1, x_2), y_2) = (\text{Ext}_1(x_1, y_1), z_2)$, where (y_1, z_2) is obtained by partitioning $\text{Ext}_2(x_2, y_2)$ into a prefix y_1 of length d_1 and a suffix z_2 of length $m_2 - d_1$.

Then for every (k_1, k_2) block source $X = (X_1, X_2)$ taking values in $\{0, 1\}^{n_1} \times \{0, 1\}^{n_2}$, it holds that $\text{Ext}'(X, U_{d_2})$ is $(\varepsilon_1 + \varepsilon_2)$ -close to $U_{m_1} \times U_{m_2 - d_1}$.

Proof. Since X_2 is a k_2 -source conditioned on any value of X_1 and Ext_2 is a (k_2, ε_2) -extractor, it follows that $(X_1, Y_1, Z_2) = (X_1, \text{Ext}_2(X_2, U_{d_2}))$ is ε_2 -close to $(X_1, U_{m_2}) = (X_1, U_{d_1}, U_{m_2 - d_1})$.

Thus, $(\text{Ext}_1(X_1, Y_1), Z_2)$ is ε_2 -close to $(\text{Ext}_1(X_1, U_{d_1}), U_{m_2 - d_1})$, which is ε_1 -close to $(U_{m_1}, U_{m_2 - d_1})$ because X_1 is a k_1 -source and Ext_1 is a (k_1, ε_1) -extractor.

By the triangle inequality, $\text{Ext}'(X, U_{d_2}) = (\text{Ext}_1(X_1, Y_1), Z_2)$ is $(\varepsilon_1 + \varepsilon_2)$ -close to $(U_{m_1}, U_{m_2 - d_1})$. \square

The benefit of this composition is that the seed length of Ext' depends only one of the extractors (namely Ext_2) rather than being the sum of the seed lengths. (If this is reminiscent of the zig-zag product, it is because they are closely related — see Section 6.3.5). Thus, we get to extract from multiple blocks at the “price of one.” Moreover, since we can take $d_1 = m_2$, which is typically much larger than d_2 , the seed length of Ext' can even be smaller than that of Ext_1 .

The lemma extends naturally to extracting from many blocks:

Lemma 6.28. For $i = 1, \dots, t$, let $\text{Ext}_i : \{0, 1\}^{n_i} \times \{0, 1\}^{d_i} \rightarrow \{0, 1\}^{m_i}$ be a (k_i, ε_i) -extractor, and suppose that $m_i \geq d_{i-1}$ for every $i = 1, \dots, t$, where we define $d_0 = 0$. Define $\text{Ext}'((x_1, \dots, x_t), y_t) = (z_1, \dots, z_t)$, where for $i = t, \dots, 1$, we inductively define (y_{i-1}, z_i) to be a partition of $\text{Ext}_i(x_i, y_i)$ into a d_{i-1} -bit prefix and a $(m_i - d_{i-1})$ -bit suffix.

Then for every (k_1, \dots, k_t) block source $X = (X_1, \dots, X_t)$ taking values in $\{0, 1\}^{n_1} \times \dots \times \{0, 1\}^{n_t}$, it holds that $\text{Ext}'(X, U_{d_t})$ is ε -close to U_m for $\varepsilon = \sum_{i=1}^t \varepsilon_i$ and $m = \sum_{i=1}^t (m_i - d_{i-1})$.

We remark that this composition preserves “strongness.” If each of the Ext_i ’s correspond to strong extractors in the sense that their seeds are prefixes of their outputs, then Ext' will also correspond to a strong extractor. If in addition $d_1 = d_2 = \dots = d_t$, then this construction can be seen as simply using the same seed to extract from all blocks.

Already with this simple composition, we can simulate **BPP** with an unpredictable-bit source (even though deterministic extraction from such sources is impossible by Proposition 6.6). As noted above, by breaking an unpredictable-bit source X with parameter δ into blocks of length ℓ , we obtain a $t \times k$ block source for $t = n/\ell$, $k = \delta'\ell$, and $\delta' = \log(1/(1 - \delta))$.

Suppose that δ is a constant. Set $\ell = (10/\delta') \log n$, so that X is a $t \times k$ block source for $k = 10 \log n$, and define $\varepsilon = n^{-2}$. Letting $\text{Ext} : \{0, 1\}^\ell \times \{0, 1\}^d \rightarrow \{0, 1\}^{d+m}$ be the (k, ε) extractor using universal hash functions (Theorem 6.18), we have:

$$\begin{aligned} d &= O(\ell) = O(\log n) & \text{and} \\ m &= k - 2 \log \frac{1}{\varepsilon} - O(1) > k/2 \end{aligned}$$

Composing Ext with itself t times as in Lemma 6.27, we obtain $\text{Ext}' : \{0, 1\}^{t\ell} \times \{0, 1\}^d \rightarrow \{0, 1\}^{d+t\cdot m}$ such that $\text{Ext}'(X, U_d)$ is ε' -close to uniform, for $\varepsilon' = 1/n$. (Specifically, $\text{Ext}'((x_1, \dots, x_t), h) = (h, h(x_1), \dots, h(x_t))$.) This tells us that Ext' essentially extracts half of the min-entropy from X , given a random seed of logarithmic length. Plugging this extractor into the construction of Proposition 6.15 gives us the following result.

Theorem 6.29. For every constant $\delta > 0$, we can simulate **BPP** with an unpredictable-bit source of parameter δ . More precisely, for every $L \in \mathbf{BPP}$ and every constant $\delta > 0$, there is a polynomial-time algorithm A and a polynomial q such that for every $w \in \{0, 1\}^*$ and every source $X \in \text{UnpredBits}_{q(|w|), \delta}$, the probability that $A(w; X)$ errs is at most $1/|w|$.

6.3.2 Reducing General Sources to Block Sources

Given the results of the previous section, a common approach to constructing extractors for general k -sources is to reduce the case of general k -sources to that of block sources.

One approach to doing this is as follows. Given a k -source X of length n , where $k = \delta n$, pick a (pseudo)random subset S of the bits of X , and let $W = X|_S$ be the bits of X in those positions. If the set S is of size ℓ , then we expect that W will have at least roughly $\delta\ell$ bits of min-entropy (with high probability over the choice of S). Moreover, W reveals at most ℓ bits of information, so if $\ell < \delta n$, intuitively there should still be min-entropy left in X . (This is justified by Lemma 6.30 below.) Thus, the pair (W, X) should be a block source. This approach can be shown to work

for appropriate ways of sampling the set S , and recursive applications of it formed the original approach to constructing good extractors (and is still useful in various contexts today). The fact mentioned above, that conditioning on a string of length ℓ reduces min-entropy by at most ℓ bits, is given by the following lemma (which is very useful when working with min-entropy) .

Lemma 6.30 (chain rule for min-entropy). If (W, X) are two jointly distributed random variables, where (W, X) is a k -source and $|\text{Supp}(W)| \leq 2^\ell$, then for every $\varepsilon > 0$, it holds that with probability at least $1 - \varepsilon$ over $w \xleftarrow{R} W$, $X|_{W=w}$ is a $(k - \ell - \log(1/\varepsilon))$ -source.

The proof of this lemma is left as an exercise (Problem 6.1).

This is referred to as the “chain rule” for min-entropy by analogy with the chain rule for Shannon entropy, which states that $H_{Sh}(X|W) = H_{Sh}(W, X) - H_{Sh}(W)$, where the conditional Shannon entropy is defined to be $H_{Sh}(X|W) = E_{w \xleftarrow{R} W} [H_{Sh}(X|_{W=w})]$. Thus, if $H_{Sh}(W, X) \geq k$ and W is of length at most ℓ , we have $H_{Sh}(X|W) \geq k - \ell$. The chain rule for min-entropy is not quite as clean; we need to assume that W has small support (rather than just small min-entropy) and we lose $\log(1/\varepsilon)$ bits of additional min-entropy. (The $\log(1/\varepsilon)$ bits of entropy can be saved by using an appropriate notion of conditional min-entropy; see Problems 6.7 and 6.8.)

Another approach, which we will follow, is based on the observation that every source of high min-entropy rate (namely, greater than $1/2$) is (close to) a block source, as shown by the lemma below. Thus, we will try to convert arbitrary sources into ones of high min-entropy rate.

Lemma 6.31. If X is an $(n - \Delta)$ -source of length n , and $X = (X_1, X_2)$ is a partition of X into blocks of lengths n_1 and n_2 , then for every $\varepsilon > 0$, (X_1, X_2) is ε -close to some $(n_1 - \Delta, n_2 - \Delta - \log(1/\varepsilon))$ block source.

Consider $\Delta = \alpha n$ for a constant $\alpha < 1/2$, and $n_1 = n_2 = n/2$. Then each block contributes min-entropy at least $(1/2 - \alpha)n$. The proofs of Lemmas 6.30 and 6.31 are left as exercises (Problem 6.1).

6.3.3 Condensers

The previous section left us with the problem of converting a general k -source into one of high min-entropy rate. We will do this via the following kind of object:

Definition 6.32. A function $\text{Con} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a $k \rightarrow_\varepsilon k'$ condenser if for every k -source X on $\{0, 1\}^n$, $\text{Con}(X, U_d)$ is ε -close to some k' -source. Con is *lossless* if $k' = k + d$.

If $k'/m > k/n$, then the condenser increases the min-entropy rate, intuitively making extraction an easier task. Indeed, condensers with $k' = m$ are simply extractors themselves.

Like extractors, it is often useful to view condensers graph-theoretically. Specifically, we think of Con as the neighbor function of a bipartite multigraph G with $N = 2^n$ left vertices, left degree $D = 2^d$, and $M = 2^m$ right vertices.

The *lossless* condensing property of Con turns out to be equivalent to G having vertex expansion close to the degree:

Proposition 6.33. Let $n, d, m \in \mathbb{N}$, $K = 2^k \in \mathbb{N}$, and $\varepsilon > 0$. A function $\text{Con} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a $k \rightarrow_\varepsilon k + d$ lossless condenser if and only if the corresponding bipartite multigraph $G = ([N], [M], E)$ of left degree D is an $(= K, (1 - \varepsilon)D)$ vertex expander.

Proof. \Rightarrow : Let $S \subset [N]$ be any set of size K . Then U_S is a k -source, so $\text{Con}(U_S, U_d)$ is ε -close to a $(k + d)$ -source. This implies that $|\text{Supp}(\text{Con}(U_S, U_d))| \geq (1 - \varepsilon) \cdot 2^{k+d}$. Noting that $\text{Supp}(\text{Con}(U_S, U_d)) = N(S)$ completes the proof.

\Leftarrow : By Lemma 6.10, it suffices to prove that for every subset $S \subset [N]$ of size K , it holds that $\text{Con}(U_S, U_d)$ is ε -close to a $(k + d)$ -source. By expansion, we know that $|N(S)| \geq (1 - \varepsilon) \cdot KD$. Since there are only KD edges leaving S , by redirecting εKD of the edges, we can ensure that they all go to KD distinct vertices in $[M]$. The uniform distribution on these KD vertices is a $(k + d)$ -source that is ε -close to $\text{Con}(U_S, U_d)$. \square

Recall that vertex expansion normally corresponds to the *disperser* property (see Proposition 6.20), which is weaker than extraction and condensing. Indeed, vertex expansion generally does not guarantee much about the distribution induced on a random neighbor of a set S , except that its support is large. However, in case the expansion is very close to the degree ($A = (1 - \varepsilon) \cdot D$), then the distribution must be nearly flat (as noted in the above proof).

By applying Proposition 6.33 to the expanders based on Parvaresh–Vardy codes (Theorem 5.35), we get the following lossless condenser:

Theorem 6.34. For every constant $\alpha > 0$, for all positive integers $n \geq k$ and all $\varepsilon > 0$, there is an explicit

$$k \rightarrow_\varepsilon k + d$$

lossless condenser $\text{Con} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = O(\log n + \log(1/\varepsilon))$ and $m = (1 + \alpha)k + O(\log(n/\varepsilon))$.

Note that setting α to be a small constant, we obtain an output min-entropy rate arbitrarily close to 1.

Problem 6.5 gives a simple extractor Ext for high min-entropy rate when the error parameter ε is constant. Applying that extractor to the output of the above condenser, we obtain extractors with a seed length of $d = O(\log n)$ that extract $\Omega(k)$ almost-uniform bits (with constant error ε) from sources of any desired min-entropy k . In the next section, we will use the above condenser to give an efficient construction of extractors for arbitrary values of ε .

We remark that the fact that having output min-entropy rate bounded away from 1 is not inherent for lossless condensers. Nonconstructively, there exist lossless condensers with output length $m = k + d + \log(1/\varepsilon) + O(1)$, and Open Problem 5.36 about expanders can be restated in the language of lossless condensers as follows:

Open Problem 6.35. Give an explicit construction of a $k \rightarrow_\varepsilon k + d$ lossless condenser $\text{Con} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = O(\log n)$, $m = k + d + O(1)$, and $\varepsilon = .01$.

If we had such a condenser, then we could get extractors that extract *all* but $O(1)$ bits of the min-entropy by then applying extractors based on spectral expanders (Theorem 6.22).

6.3.4 The Extractor

In this section, we will use the ideas outlined in the previous section — namely condensing and block-source extraction — to construct an extractor that is optimal up to constant factors.

Theorem 6.36. For all positive integers $n \geq k$ and all $\varepsilon > 0$, there is an explicit (k, ε) extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $m \geq k/2$ and $d = O(\log(n/\varepsilon))$.

We will use the following building block, constructed in Problem 6.9.

Lemma 6.37. For every constant $t > 0$ and all positive integers $n \geq k$ and all $\varepsilon > 0$, there is an explicit (k, ε) extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $m \geq k/2$ and $d = k/t + O(\log(n/\varepsilon))$.

The point is that this extractor has a seed length that is an arbitrarily large constant factor (approximately $t/2$) smaller than its output length. Thus, if we use it as Ext_2 in the block-source extraction of Lemma 6.27, the resulting seed length will be smaller than that of Ext_1 by an arbitrarily large constant factor. (The seed length of the composed extractor Ext' in Lemma 6.27 is the same of that as Ext_2 , which will be a constant factor smaller than its output length m_2 , which we can take to be equal to the seed length d_1 of Ext_1 .)

Overview of the Construction. Note that for small min-entropies k , namely $k = O(\log(n/\varepsilon))$, the extractor we want is already given by Lemma 6.37 with seed length d smaller than the output length m by any constant factor. (If we allow $d \geq m$, then extraction is trivial — just output the seed.) Thus, our goal will be to recursively construct extractors for large min-entropies using extractors for smaller min-entropies. Of course, if $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k_0, ε) extractor, say with $m = k_0/2$, then it is also a (k, ε) extractor for every $k \geq k_0$. The problem is that the output length is only $k_0/2$ rather than $k/2$. Thus, we need to increase the output length. This can be achieved by simply applying extractors for smaller min-entropies several times:

Lemma 6.38. Suppose $\text{Ext}_1 : \{0, 1\}^n \times \{0, 1\}^{d_1} \rightarrow \{0, 1\}^{m_1}$ is a (k_1, ε_1) extractor and $\text{Ext}_2 : \{0, 1\}^n \times \{0, 1\}^{d_2} \rightarrow \{0, 1\}^{m_2}$ is a (k_2, ε_2) extractor for $k_2 = k_1 - m_1 - \log(1/\varepsilon_3)$. Then $\text{Ext}' : \{0, 1\}^n \times \{0, 1\}^{d_1+d_2} \rightarrow \{0, 1\}^{m_1+m_2}$ defined by $\text{Ext}'(x, (y_1, y_2)) = (\text{Ext}_1(x, y_1), \text{Ext}_2(x, y_2))$ is a $(k_1, \varepsilon_1 + \varepsilon_2 + \varepsilon_3)$ extractor.

The proof of this lemma follows from Lemma 6.30. After conditioning a k_1 -source X on $W = \text{Ext}_1(X, U_{d_1})$, X still has min-entropy at least $k_1 - m_1 - \log(1/\varepsilon_3) = k_2$ (except with probability ε_3), and thus $\text{Ext}_2(X, U_{d_2})$ can extract an additional m_2 almost-uniform bits.

To see how we might apply this, consider setting $k_1 = .8k$ and $m_1 = k_1/2$, $\varepsilon_1 = \varepsilon_2 = \varepsilon_3 = \varepsilon \geq 2^{-.1k}$, $k_2 = k_1 - m_1 - \log(1/\varepsilon_3) \in [.3k, .4k]$, and $m_2 = k_2/2$. Then we obtain a $(k, 3\varepsilon)$ extractor Ext' with output length $m = m_1 + m_2 > k/2$ from two extractors for min-entropies k_1, k_2 that are smaller than k by a constant factor, and we can hope to construct the latter two extractors recursively via the same construction.

Now, however, the problem is that the seed length grows by a constant factor in each level of recursion (e.g. if $d_1 = d_2 = d$ in Lemma 6.38, we get seed length $2d$ rather than d). Fortunately,

block source extraction using the extractor of Lemma 6.37 gives us a method to reduce the seed length by a constant factor. (The seed length of the composed extractor Ext' in Lemma 6.27 is the same of that as Ext_2 , which will be a constant factor smaller than its output length m_2 , which we can take to be equal to the seed length d_1 of Ext_1 . Thus, the seed length of Ext' will be a constant factor smaller than that of Ext_1 .) In order to apply block source extraction, we first need to convert our source to a block source; by Lemma 6.31, we can do this by using the condenser of Theorem 6.34 to make its entropy rate close to 1.

One remaining issue is that the error ε still grows by a constant factor in each level of recursion. However, we can start with polynomially small error at the base of the recursion and there are only logarithmically many levels of recursion, so we can afford this blow-up.

We now proceed with the proof details. It will be notationally convenient to do the steps in the reverse order from the description above — first we will reduce the seed length by a constant factor via block-source extraction, and then apply Lemma 6.38 to increase the output length.

Proof of Theorem 6.36. Fix $n \in \mathbb{N}$ and $\varepsilon_0 > 0$. Set $d = c \log(n/\varepsilon_0)$ for an error parameter ε_0 and a sufficiently large constant c to be determined in the proof below. (To avoid ambiguity, we will keep the dependence on c explicit throughout the proof, and all big-Oh notation hides universal constants independent of c .) For $k \in [0, n]$, let $i(k)$ be the smallest nonnegative integer i such that $k \leq 2^i \cdot 8d$. This will be the level of recursion in which we handle min-entropy k ; note that $i(k) \leq \log k \leq \log n$.

For every $k \in [0, n]$, we will construct an explicit $\text{Ext}_k : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^{k/2}$ that is a $(k, \varepsilon_{i(k)})$ extractor, for an appropriate sequence $\varepsilon_0 \leq \varepsilon_1 \leq \varepsilon_2 \cdots$. Note that we require the seed length to remain d and the fraction of min-entropy extracted to remain $1/2$ for all values of k . The construction will be by induction on $i(k)$.

Base Case: $i(k) = 0$, i.e. $k \leq 8d$. The construction of Ext_k follows from Lemma 6.37, setting $t = 9$ and taking c to be a sufficiently large constant.

Inductive Case: We construct Ext_k for $i(k) \geq 1$ from extractors $\text{Ext}_{k'}$ with $i(k') < i(k)$ as follows. Given a k -source X of length n , Ext_k works as follows.

- (1) We apply the condenser of Theorem 6.34 to convert X into a source X' that is ε_0 -close to a k -source of length $(9/8)k + O(\log(n/\varepsilon_0))$. This requires a seed of length $O(\log(n/\varepsilon_0))$.
- (2) We divide X' into two equal-sized halves (X_1, X_2) . By Lemma 6.31, (X_1, X_2) is $2\varepsilon_0$ -close to a $2 \times k'$ block source for

$$k' = k/2 - k/8 - O(\log(n/\varepsilon_0)) .$$

Note that $i(k') < i(k)$. Since $i(k) \geq 1$, we also have $k' \geq 3d - O(\log(n/\varepsilon_0)) \geq 2d$, for a sufficiently large choice of the constant c .

- (3) Now we apply block-source extraction as in Lemma 6.27. We take Ext_2 to be a $(2d, \varepsilon_0)$ extractor from Lemma 6.37 with parameter $t = 16$, which will give us $m_2 = d$ output bits using a seed of length $d_2 = (2d)/16 + O(\log(n/\varepsilon_0))$. For Ext_1 , we use our recursively constructed $\text{Ext}_{k'}$, which has seed length d , error $\varepsilon_{i(k')}$, and output length $k'/2 \geq k/6$ (where the latter inequality holds for a sufficiently large choice of the constant c , because $k > 8d > 8c \log(1/\varepsilon)$).

All in all, our extractor so far has seed length at most $d/8 + O(\log(n/\varepsilon_0))$, error at most $\varepsilon_{i(k)-1} + O(\varepsilon_0)$, and output length at least $k/6$. This would be sufficient for our induction except that the output length is only $k/6$ rather than $k/2$. We remedy this by applying Lemma 6.38.

With one application of the extractor above, we extract at least $m_1 = k/6$ bits of the source min-entropy. Then with another application of the extractor above for min-entropy threshold $k_2 = k - m_1 - \log(1/\varepsilon_0) = 5k/6 - \log(1/\varepsilon_0)$, by Lemma 6.38, we extract another $(5k/6 - \log(1/\varepsilon_0))/6$ bits and so on. After four applications, we have extracted all but $(5/6)^4 \cdot k + O(\log(1/\varepsilon_0)) \leq k/2$ bits of the min-entropy. Our seed length is then $4 \cdot (d/8 + O(\log(n/\varepsilon_0))) \leq d$ and the total error is $\varepsilon_{i(k)} = O(\varepsilon_{i(k)-1})$.

Solving the recurrence for the error, we get $\varepsilon_i = 2^{O(i)} \cdot \varepsilon_0 \leq \text{poly}(n) \cdot \varepsilon_0$, so we can obtain error ε by setting $\varepsilon_0 = \varepsilon/\text{poly}(n)$. As far as explicitness, we note that computing Ext_k consists of four evaluations of our condenser from Theorem 6.34, four evaluations of $\text{Ext}_{k'}$ for values of k' such that $i(k') < (i(k) - 1)$, four evaluations of the explicit extractor from Lemma 6.37, and simple string manipulations that can be done in time $\text{poly}(n, d)$. Thus, the total computation time is at most $4^{i(k)} \cdot \text{poly}(n, d) = \text{poly}(n, d)$. \square

Repeatedly applying Lemma 6.38 using extractors from Theorem 6.36, we can extract any constant fraction of the min-entropy using a logarithmic seed length, and all the min-entropy using a polylogarithmic seed length.

Corollary 6.39. The following holds for every constant $\alpha > 0$. For every $n \in \mathbb{N}$, $k \in [0, n]$, and $\varepsilon > 0$, there is an explicit (k, ε) extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $m \geq (1 - \alpha)k$ and $d = O(\log(n/\varepsilon))$.

Corollary 6.40. For every $n \in \mathbb{N}$, $k \in [0, n]$, and $\varepsilon > 0$, there is an explicit (k, ε) extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $m = k - O(\log(1/\varepsilon))$ and $d = O(\log k \cdot \log(n/\varepsilon))$.

We remark that the above construction can be modified to yield *strong* extractors achieving the same output lengths as above (so the entropy of the seed need not be lost in Corollary 6.40).

A summary of the extractor parameters we have seen is in Table 6.2.

Method	Seed Length d	Output Length m
Optimal and Nonconstructive	$\log(n - k) + O(1)$	$k + d - O(1)$
Necessary for BPP Simulation	$O(\log n)$	$k^{\Omega(1)}$
Spectral Expanders	$O(n - k)$	n
Pairwise Independent Hashing	$O(n)$	$k + d - O(1)$
Corollary 6.39	$O(\log n)$	$(1 - \gamma)k$, any constant $\gamma > 0$
Corollary 6.40	$O(\log^2 n)$	$k - O(1)$

Table 6.2 Parameters for some constructions of $(k, .01)$ extractors.

While Theorem 6.36 and Corollary 6.39 give extractors that are optimal up to constant factors in both the seed length and output length, it remains an important open problem to get one or both

of these to be optimal to within an *additive* constants while keeping the other optimal to within a constant factor.

Open Problem 6.41. Give an explicit construction of $(k, .01)$ extractors $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with seed length $d = O(\log n)$ and output length $m = k + d - O(1)$.

By using the condenser of Theorem 6.34, it suffices to achieve the above for high min-entropy rate, e.g. $k = .99n$. Alternatively, a better condenser construction would also resolve the problem. (See Open Problem 6.35.) We note that there is a recent construction of extractors with seed length $d = O(\log n)$ and output length $m = (1 - 1/\text{polylog}(n))k$ (improving the output length of $m = \Omega(k)$ of Corollary 6.39, in the case of constant or slightly subconstant error).

Open Problem 6.42. Give an explicit construction of $(k, .01)$ extractors $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with seed length $d = \log n + O(1)$ and $m = \Omega(k)$ (or even $m = k^{\Omega(1)}$).

One of the reasons that these open problems are significant is that, in many applications of extractors, the resulting complexity depends *exponentially* on the seed length d and/or the entropy loss $k + d - m$. (An example is the simulation of **BPP** with weak random sources given by Proposition 6.15.) Thus, additive constants in these parameters corresponds to constant multiplicative factors in complexity.

Another open problem is more aesthetic in nature. The construction of Theorem 6.36 makes use of the condenser of Theorem 6.36, the Leftover Hash Lemma (Theorem 6.18) and the composition techniques of Lemmas 6.27 and Lemma 6.38 in a somewhat complex recursion. It is of interest to have a construction that is more direct. In addition to the aesthetic appeal, such a construction would likely be more practical to implement and provide more insight into extractors. In Chapter 7, we will see a very direct construction based on a connection between extractors and pseudorandom generators, but its parameters will be somewhat worse than Theorem 6.36. Thus the following remains open:

Open Problem 6.43. Give a “direct” construction of (k, ε) extractors $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with seed length $d = O(\log(n/\varepsilon))$ and $m = \Omega(k)$.

6.3.5 Block-Source Extraction vs. the Zig-Zag Product

To further highlight the connections between extractors and expanders, here we describe how the block-source extraction method of Section 6.3.1 (which we used in the main extractor construction of Theorem 6.36) is closely related to the zig-zag graph product of Section 4.3.2.3 (which we used in the expander construction of Theorem 4.39).

Recall the block-source extraction method of Lemma 6.27: We define $\text{Ext}' : \{0, 1\}^{n_1+n_2} \times \{0, 1\}^{d_2} \rightarrow \{0, 1\}^{m_1}$ by $\text{Ext}'((x_1, x_2), y_2) = \text{Ext}_1(x_1, \text{Ext}_2(x_2, y_2))$. (Here we consider the special case that $m_2 = d_1$.)

Viewing the extractors as bipartite graphs, the left-vertex set is $[N_1] \times [N_2]$ and the left-degree is D_2 . A random step from a vertex $(x_1, x_2) \in [N_1] \times [N_2]$ corresponds to taking a random step

from x_2 in G_2 to obtain a right-hand vertex $y_1 \in \{0, 1\}^{m_2}$, which we view as an edge label y for G_1 . We then move to the y 'th neighbor of x_1 .

This is just like the first two steps of the zig-zag graph product. Why do we need a third step in the zig-zag product? It is because of the slightly different goals in the two settings. In a (spectral) expander, we consider an arbitrary initial distribution that does not have too much (Rényi) entropy, and need to add entropy to it. In a block-source extractor, our initial distribution is constrained to be a block source (so both blocks have a certain amount of min-entropy), and our goal is to produce an almost-uniform output (even if we end up with less bits than the initial entropy).

Thus, in the zig-zag setting, we must consider the following extreme cases (that are ruled out for block sources):

- The second block has no entropy given the first. Here, the step using G_2 will add entropy, but not enough to make y_1 close to uniform. Thus, we have no guarantees on the behavior of the G_1 -step, and we may lose entropy with it. For this reason, we keep track of the edge used in the G_1 -step — that is, we remember b_1 such that x_1 is the b_1 'th neighbor of $z_1 = \text{Ext}(x_1, y_1)$. This ensures that the (edge-rotation) mapping $(x_1, y_1) \mapsto (z_1, b_1)$ is a permutation and does not lose any entropy. We can think of b_1 as a “buffer” that retains any extra entropy in (x_1, y_1) that did not get extracted into z_1 . So a natural idea is to just do block source extraction, but output (z_1, b_1) rather than just z_1 . However, this runs into trouble with the next case.
- The first block has no entropy but the second block is completely uniform given the first. In this case, the G_2 step cannot add any entropy and the G_1 step does not add any entropy because it is a permutation. However, the G_1 step transfers entropy into z_1 . So if we add another expander-step from b_1 at the end, we can argue that it will add entropy. This gives rise to the 3-step definition of the zig-zag product.

While we analyzed the zig-zag product with respect to spectral expansion (i.e. Rényi entropy), it is also possible to analyze it in terms of a condenser-like definition (i.e. outputting distributions ε -close to having some min-entropy). It turns out that a variant of the zig-zag product for condensers leads to a construction of constant-degree bipartite expanders with expansion $(1 - \varepsilon) \cdot D$ for the balanced ($M = N$) or nearly balanced (e.g. $M = \Omega(N)$) case. However, as mentioned in Open Problems 4.43, 4.44, and 5.36, there are still several significant open problems concerning the explicit construction of expanders with vertex expansion close to the degree, achieving expansion $D - O(1)$ in the non-bipartite case, and achieving a near-optimal number of right-hand vertices.

6.4 Exercises

Problem 6.1 (Min-entropy and Statistical Difference). — (1) Prove that for every two random variables X and Y ,

$$\Delta(X, Y) = \max_f |\mathbb{E}[f(X)] - \mathbb{E}[f(Y)]| = \frac{1}{2} \cdot \|X - Y\|_1,$$

where the maximum is over all $[0, 1]$ -valued functions f . (Hint: first identify the functions f that maximize $|\mathbb{E}[f(X)] - \mathbb{E}[f(Y)]|$.)

- (2) Suppose that (W, X) are jointly distributed random variables where (W, X) is a k -source and $|\text{Supp}(W)| \leq 2^\ell$. Show that for every $\varepsilon > 0$, with probability at least $1 - \varepsilon$ over $w \stackrel{R}{\leftarrow} W$, we have $X|_{W=w}$ is a $(k - \ell - \log(1/\varepsilon))$ -source.
- (3) Suppose that X is an $(n - \Delta)$ -source taking values in $\{0, 1\}^n$, and we let X_1 consist of the first n_1 bits of X and X_2 the remaining $n_2 = n - n_1$ bits. Show that for every $\varepsilon > 0$, (X_1, X_2) is ε -close to some $(n_1 - \Delta, n_2 - \Delta - \log(1/\varepsilon))$ block source.
-

Problem 6.2 (Simulating Randomized Algorithms with Weak Sources). (1)

Let $A(w; r)$ be a randomized algorithm for computing a function f using m random bits such that $A(w; U_m)$ has error probability at most $1/3$ (i.e. for every w , $\Pr[A(w; U_m) \neq f(w)] \leq 1/3$) and let $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ be a $(k, 1/7)$ -extractor. Define $A'(w; x) = \text{maj}_{y \in \{0, 1\}^d} \{A(w; \text{Ext}(x; y))\}$ (breaking ties arbitrarily). Show that for every $(k + t)$ -source X , $A'(w; X)$ has error probability at most 2^{-t} .

- (2) Let $A(w; r)$ be a randomized algorithm for deciding a language L using m random bits such that $A(w; U_m)$ has one-sided error probability at most $1/2$ (i.e. if $w \in L$, then $\Pr[A(w; U_m) = 1] \geq 1/2$ and if $w \notin L$, then $\Pr[A(w; U_m) = 1] = 0$) and let $\text{Disp} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ be a $(k, 1/3)$ -disperser. Show how to decide L with one-sided error at most 2^{-t} given a single sample from a $(k + t)$ -source X (with no other randomness) and running A and Disp each 2^d times.
-

Problem 6.3 (Almost-Universal Hashing). A family \mathcal{H} of functions mapping domain $[N]$ to $[M]$ is said to have *collision probability* at most δ if for every $x_1 \neq x_2 \in [N]$, we have

$$\Pr_{H \stackrel{R}{\leftarrow} \mathcal{H}} [H(x_1) = H(x_2)] \leq \delta.$$

\mathcal{H} has is ε -almost universal if it has collision probability at most $(1 + \varepsilon)/M$. (Note that this is a relaxation of the notion of ε -almost pairwise independence from Problem 3.4.)

- (1) Show that if a family $\mathcal{H} = \{h : [N] \rightarrow [M]\}$ is ε^2 -almost universal, $\text{Ext}(x, h) \stackrel{\text{def}}{=} h(x)$ is a (k, ε) strong extractor for $k = m + 2 \log(1/\varepsilon) + O(1)$, where $m = \log M$.
- (2) Use Problem 3.4 to deduce that for every $n \in \mathbb{N}$, $k \leq n$, and $\varepsilon > 0$, there is a (k, ε) strong extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = O(k + \log(n/\varepsilon))$ and $m = k - 2 \log(1/\varepsilon) - O(1)$.
- (3) Given a family \mathcal{H} of functions mapping $[N]$ to $[M]$, we can obtain a code $\text{Enc} : [N] \rightarrow [M]^{|\mathcal{H}|}$ by $\text{Enc}(x)_h = h(x)$, and conversely. Show that \mathcal{H} has collision probability at most δ iff Enc has minimum distance at least $1 - \delta$.
- (4) Use the above connections and the list-decoding view of extractors (Proposition 6.23) to prove the Johnson Bound for small alphabets: if a code $\text{Enc} : [N] \rightarrow [M]^{\hat{n}}$ has minimum distance at least $1 - 1/M - \gamma/M$, then it is $(1 - 1/M - \sqrt{\gamma}, O(M/\gamma))$ list-decodable.
-

Problem 6.4 (Rényi extractors). Call a function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ a (k, ε) Rényi extractor if for every source X on $\{0, 1\}^n$ of Rényi entropy at least k , it holds that $\text{Ext}(X, U_d)$ has Rényi entropy at least $m - \varepsilon$.

- (1) Prove that a (k, ε) Rényi extractor is also a $(k, O(\sqrt{\varepsilon}))$ extractor.
 - (2) Show that for every $n, k, m \in \mathbb{N}$ with $m \leq n$ and $\varepsilon > 0$, there exists a (k, ε) Rényi extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = O(\min\{n - k + \log(1/\varepsilon), m/2 + \log(n/\varepsilon)\})$.
 - (3) Show that if $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a $(k, 1)$ Rényi extractor, then $d \geq \min\{n - k, m/2\} - O(1)$. (Hint: consider a k -source that is uniform over $\{x : \exists y \text{Ext}(x, y) \in T\}$ for an appropriately chosen set T .)
-

Problem 6.5 (Extractors vs. Samplers). Use the connection between extractors and averaging samplers to do the following:

- (1) Prove that for all constants $\varepsilon, \alpha > 0$, there is a constant $\beta < 1$ such that for all n , there is an explicit $(\beta n, \varepsilon)$ extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d \leq \log n$ and $m \geq (1 - \alpha)n$.
 - (2) Prove that for every $m \in \mathbb{N}$, $\varepsilon, \delta > 0$, there exists a (nonconstructive) (δ, ε) averaging sampler $\text{Samp} : \{0, 1\}^n \rightarrow (\{0, 1\}^m)^t$ using $n = m + 2 \log(1/\varepsilon) + \log(1/\delta) + O(1)$ random bits and $t = O((1/\varepsilon^2) \cdot \log(1/\delta))$ samples.
 - (3) Suppose we are given a constant-error **BPP** algorithm that uses $r = r(n)$ random bits on inputs of length n . Show how, using the explicit extractor of Theorem 6.36, we can reduce its error probability to $2^{-\ell}$ using $O(r) + \ell$ random bits, for any polynomial $\ell = \ell(n)$. (Note that this improves the $r + O(\ell)$ given by expander walks for $\ell \gg r$.) Conclude that every problem in **BPP** has a randomized polynomial-time algorithm that only errs for $2^{q^{0.01}}$ choices of its $q = q(n)$ random bits.
-

Problem 6.6 (Extracting from Unpredictable-Bit Sources). (1) Let X be a source taking values in $\{0, 1\}^n$ such that for all x, y , $\Pr[X = x] / \Pr[X = y] \leq (1 - \delta) / \delta$. Show that $X \in \text{UnpredBits}_{n, \delta}$.

- (2) Prove that for every function $\text{Ext} : \{0, 1\}^n \rightarrow \{0, 1\}$ and every $\delta > 0$, there exists a source $X \in \text{UnpredBits}_{n, \delta}$ with parameter δ such that $\Pr[\text{Ext}(X) = 1] \leq \delta$ or $\Pr[\text{Ext}(X) = 0] \geq 1 - \delta$. (Hint: for $b \in \{0, 1\}$, consider X that is uniform on $\text{Ext}^{-1}(b)$ with probability $1 - \delta$ and is uniform on $\text{Ext}^{-1}(\bar{b})$ with probability δ .)
 - (3) (*) Show how to extract from sources in $\text{UnpredBits}_{n, \delta}$ using a seeded extractor with a seed of *constant* length. That is, the seed length should not depend on the length n of the source, but only on the bias parameter δ and the statistical difference ε from uniform desired. The number of bits extracted should be $\Omega(\delta n)$.
-

Problem 6.7 (Average Min-Entropy). While there is no single “correct” definition of conditional min-entropy, in this problem you will explore one definition that is useful and has nice properties. For two jointly distributed random variables (X, Y) define

$$H_\infty(Y|X) = \log \left(\frac{1}{\mathbb{E}_{x \leftarrow X} [\max_y \Pr[Y = y|X = x]]} \right).$$

- (1) Prove that $H_\infty(Y|X) = \log(1/\max_P \Pr[P(X) = Y])$, where the maximum is over functions P from the domain of X to the domain of Y . So average min-entropy captures the unpredictability of Y from X .
 - (2) Show that if $H_\infty(Y|X) \geq k$, then with probability at least $1 - \varepsilon$ over $x \leftarrow X$, we have $H_\infty(Y|_{X=x}) \geq k - \log(1/\varepsilon)$.
 - (3) Show that $H_\infty(Y) \geq H_\infty(Y|X) \geq H_\infty(X, Y) - \log |\text{Supp}(X)|$.
 - (4) Use the previous two items to deduce Problem 6.1, Part 2.
 - (5) Give an example showing that $H_\infty(Y|X)$ can be much smaller than $H_\infty(Y) - H_\infty(X)$. Specifically, construct n -bit random variables where $H_\infty(Y) = n$ but $H_\infty(Y|X)$ and $H_\infty(X)$ are both $O(1)$.
-

Problem 6.8 (Average Min-Entropy Extractors). A function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε) average min-entropy extractor if for every two jointly distributed random variables (X, Y) such that X takes values in $\{0, 1\}^n$ and $H_\infty(X|Y) \geq k$, we have $(\text{Ext}(X, U_d), Y)$ is ε -close to (U_m, Y) , where U_d and U_m are taken to be independent of (X, Y) . By Problem 6.7, Part 2, it follows that if Ext is a (k, ε) extractor, then it is a $(k + \log(1/\varepsilon), 2\varepsilon)$ average min-entropy extractor. Below, you will show how to avoid the $\log(1/\varepsilon)$ entropy loss from this reduction.

- (1) (*) Show that if $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε) -extractor with $k \leq n - 1$, then for every $t > 0$, Ext is a $(k - t, 2^{t+1} \cdot \varepsilon)$ -extractor. (Hint: for a statistical test $T \subseteq \{0, 1\}^m$, relate the $(k - t)$ -sources X maximizing the distinguishing advantage $|\Pr[\text{Ext}(X, U_d) \in T] - \Pr[U_m \in T]|$ to the k -sources maximizing the distinguishing advantage.)
 - (2) Show that if $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε) -extractor, then Ext is a $(k, 3\varepsilon)$ average min-entropy extractor.
 - (3) Use these results and Problem 6.3 to improve Corollary 6.40, and construct, for every $n \in \mathbb{N}$, $k \leq n$, and $\varepsilon > 0$, an explicit (k, ε) extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with seed length $d = O((\log k) \cdot (\log(n/\varepsilon)))$ and output length $m = k + d - 2 \log(1/\varepsilon) - O(1)$.
-

Problem 6.9 (The Building-Block Extractor). Prove Lemma 6.37: Show that for every constant $t > 0$ and all positive integers $n \geq k$ and all $\varepsilon > 0$, there is an explicit (k, ε) -extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $m \geq k/2$ and $d = k/t + O(\log(n/\varepsilon))$. (Hint: convert the source into a block source with blocks of length $k/O(t) + O(\log(n/\varepsilon))$.)

Problem 6.10 (Encryption and Deterministic Extraction). A (one-time) encryption scheme with key length n and message length m consists of an encryption function $\text{Enc}: \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^\ell$ and a decryption function $\text{Dec}: \{0, 1\}^\ell \times \{0, 1\}^m \rightarrow \{0, 1\}^m$ such that $\text{Dec}(k, \text{Enc}(k, u)) = u$ for every $k \in \{0, 1\}^n$ and $u \in \{0, 1\}^m$. Let K be a random variable taking values in $\{0, 1\}^n$. We say that (Enc, Dec) is (statistically) ε -secure with respect to K if for every two messages $u, v \in \{0, 1\}^m$, we have $\Delta(\text{Enc}(K, u), \text{Enc}(K, v)) \leq \varepsilon$. For example, the one-time pad, where $n = m = \ell$ and $\text{Enc}(k, u) = k \oplus u = \text{Dec}(k, u)$ is 0-secure (aka perfectly secure) with respect to the uniform distribution $K = U_m$. For a class \mathcal{C} of sources on $\{0, 1\}^n$, we say that the encryption scheme (Enc, Dec) is ε -secure with respect to \mathcal{C} if Enc is ε -secure with respect to every $K \in \mathcal{C}$.

- (1) Show that if there exists a deterministic ε -extractor $\text{Ext}: \{0, 1\}^n \rightarrow \{0, 1\}^m$ for \mathcal{C} , then there exists a 2ε -secure encryption scheme with respect to \mathcal{C} .
- (2) Conversely, use the following steps to show that if there exists an ε -secure encryption scheme (Enc, Dec) with respect to \mathcal{C} , where $\text{Enc}: \{0, 1\}^n \times \{0, 1\}^m \rightarrow \{0, 1\}^\ell$, then there exists a deterministic 2ε -extractor $\text{Ext}: \{0, 1\}^n \rightarrow \{0, 1\}^{m-2\log(1/\varepsilon)-O(1)}$ for \mathcal{C} , provided $m \geq \log n + 2\log(1/\varepsilon) + O(1)$.
 - (a) For each fixed key $k \in \{0, 1\}^n$, define a source X_k on $\{0, 1\}^\ell$ by $X_k = \text{Enc}(k, U_m)$, and let \mathcal{C}' be the class of all these sources (i.e., $\mathcal{C}' = \{X_k : k \in \{0, 1\}^n\}$). Show that there exists a deterministic ε -extractor $\text{Ext}': \{0, 1\}^\ell \rightarrow \{0, 1\}^{m-2\log(1/\varepsilon)-O(1)}$ for \mathcal{C}' , provided $m \geq \log n + 2\log(1/\varepsilon) + O(1)$.
 - (b) Show that if Ext' is a deterministic ε -extractor for \mathcal{C}' and Enc is ε -secure with respect to \mathcal{C} , then $\text{Ext}(k) = \text{Ext}'(\text{Enc}(k, 0^m))$ is a deterministic 2ε -extractor for \mathcal{C} .

Thus, a class of sources can be used for secure encryption iff it is deterministically extractable.

Problem 6.11 (Extracting from Symbol-Fixing Sources*). A generalization of a bit-fixing source is a *symbol-fixing source* X taking values in Σ^n for some alphabet Σ , where subset of the coordinates of X are fixed and the rest are uniformly distributed and independent elements of Σ . For $\Sigma = \{0, 1, 2\}$ and $k \in [0, n]$, give an explicit ε -extractor $\text{Ext}: \Sigma^n \rightarrow \{0, 1\}^m$ for the class of symbol-fixing sources on Σ^n with min-entropy at least k , with $m = \Omega(k)$ and $\varepsilon = 2^{-\Omega(k)}$. (Hint: use a random walk on a consistently labelled 3-regular expander graph.)

6.5 Chapter Notes and References

Surveys on randomness extractors and their applications are given by Nisan and Ta-Shma [NT] and Shaltiel [Sha1, Sha3].

The study of randomness extraction goes back to von Neumann [von], who gave the deterministic extractor described in the text for IID-bit sources. (See [Eli2, Per] for improvements in the extraction rate.) The study of extraction was renewed by Blum [Blu], who showed how to extend

von Neumann’s method to extract perfectly uniform bits from a source generated by a finite-state Markov chain. Santha and Vazirani [SV] proposed the model of unpredictable-bit sources, suggested that an extractor need only produce an output that is statistically close to uniform (rather than perfectly uniform), and proved that there is no deterministic extractor for a single unpredictable-bit source even under this relaxation (Proposition 6.6; the proof we give in Problem 6.6 is from [RVW3]). Vazirani and Vazirani [Vaz1] showed that nevertheless every problem in **BPP** can be solved in polynomial time given a single unpredictable-bit source (Theorem 6.29.) Chor and Goldreich [CG1] generalized this result to block sources, in the process introducing min-entropy to the literature on randomness extraction. Cohen and Wigderson [CW1] showed how to simulate **BPP** with any source of sufficiently high min-entropy rate. This was strengthened to any constant entropy rate by Zuckerman [Zuc1], and polynomially small entropy rate in [SSZ, ACRT].

The notion of seeded extractor we focus on in this chapter (Definitions 6.13,6.16) was introduced by Nisan and Zuckerman [NZ], who also gave the first construction of extractors with polylogarithmic seed length (building on [Zuc1]). A form of the Leftover Hash Lemma (for flat distributions, and where the quality of the output distribution is measured with entropy rather than statistical difference) was first proven by Bennett, Brassard, and Robert [BBR]. The version in Theorem 6.18 is due to Impagliazzo, Levin, and Luby [HILL], and the proof we give is due to Rackoff [IZ]. The connection between the Leftover Hash Lemma and the Johnson Bound in Problem 3.4 is due to Ta-Shma and Zuckerman [TZ], with the equivalence between almost-universal hashing and minimum distance being from [BJKS]. The extractor parameters of Problem 3.4 were first obtained in [Gol1, SZ].

The study of deterministic extractors has also remained active, motivated by applications in cryptography and other settings where enumerating seeds does not work. Bit-fixing sources, first studied in [Vaz3, BBR, CGH⁺], received particular attention for their applications to maintaining security when an adversary learns some bits of a cryptographic secret key [CDH⁺]. (There has been work on protecting against even more severe forms of leakage, surveyed in [Gol].) The deterministic extractor for symbol-fixing sources in Problem 6.11 is from [KZ]. Problem 6.10 (encryption requires deterministic extraction) is due to Bosley and Dodis [BD]. For a discussion of deterministic extraction for other models of sources, see Section 8.2.

The relevance of bipartite expanders to simulating randomized algorithms with weak sources is due to Cohen and Wigderson [CW1], who studied dispersers and generalizations of them (as graphs rather than functions). The expander-based extractor of Theorem 6.22 is due to Goldreich and Wigderson [GW]. The connection between extractors and list-decodable codes emerged in the work of Trevisan [Tre1], and the list-decoding view of extractors (Proposition 6.23) was crystallized by Ta-Shma and Zuckerman [TZ]. The equivalence between extractors and averaging samplers (Corollary 6.24) is due to Zuckerman [Zuc2] (building on previous connections between other types of samplers and expanders/dispersers [Sip2, CW1]).

The notion of entropy was introduced by Shannon in his seminal paper [Sha3] that gave rise to the field of information theory. The many properties of Shannon entropy and related quantities and their applications to communications theory are covered in the textbook [CT]. There, the fact that Shannon entropy converges to min-entropy (up to a vanishing statistical distance) when we take independent samples is called the Asymptotic Equipartition Property. Rényi entropy (indeed the entropies H_α for all $\alpha \in (0, \infty)$) were introduced in [Rén]. The fact that every k -source is a convex combination of flat k -sources (Lemma 6.10) is implicit in [CG1]. Our proof is due to Kalai [Kal].

The optimal nonconstructive bounds for extractor parameters (Theorem 6.14) were identified by Radhakrishnan and Ta-Shma [RT], who proved a matching lower bound on seed length and upper bound on output length.

The method for extracting from block sources given by Lemmas 6.27 and 6.28 was developed over a series of works [VV, CG1, Zuc1]; the form we use is from [NZ]. The first method described in Section 6.3.2 for converting general weak sources to block sources is from the original papers of Nisan and Zuckerman [Zuc1, NZ] (see [Vad] for a tighter analysis). The observation that high min-entropy sources can be partitioned into block sources, and the benefits of this for constructing extractors, is due to Goldreich and Wigderson [GW]. Lossless condensers were first introduced by Raz and Reingold [RR] (formulated in a slightly different manner), and the general definition of condensers is from Reingold, Shaltiel, and Wigderson [RSW]. The equivalence of lossless condensing and vertex expansion close to the degree is from Ta-Shma, Umans, and Zuckerman [TUZ]. The extractor construction of Section 6.3.4 is due to Guruswami, Umans, and Vadhan [GUV]. (Earlier papers had results similar to Theorem 6.36 with some restrictions on the parameters, such as $k = \Omega(n)$ [Zuc2] and $1/\varepsilon$ subpolynomial [LRVW].) Recently, Dvir, Kopparty, Saraf, and Sudan [DKSS] constructed extractors that have logarithmic seed length and extract a $1 - o(1)$ fraction of the min-entropy (for constant or slightly subconstant error).

The zig-zag product for extractors and condensers described in Section 6.3.5 is from [RVW1, CRVW].

Problem 6.2 (achieving exponentially small error when simulating **BPP** with a weak source) is from [Zuc1].

7

Pseudorandom Generators

7.1 Motivation and Definition

In the previous chapters, we have seen a number of interesting derandomization results:

- Derandomizing specific algorithms, such as the ones for MAXCUT and UNDIRECTED S-T CONNECTIVITY;
- Giving explicit (efficient, deterministic) constructions of various pseudorandom objects, such as expanders, extractors, and list-decodable codes, as well as showing various relations between them;
- Reducing the randomness needed for certain tasks, such as sampling and amplifying the success probability of randomized algorithm; and
- Simulating **BPP** with any weak random source.

However, all of these still fall short of answering our original motivating question, of whether *every* randomized algorithm can be efficiently derandomized. That is, does **BPP** = **P**?

As we have seen, one way to resolve this question in the positive is to use the following two-step process: First show that the number of random bits for any **BPP** algorithm can be reduced from $\text{poly}(n)$ to $O(\log n)$, and then eliminate the randomness entirely by enumeration.

Thus, we would like to have a function G that stretches a seed of $d = O(\log n)$ truly random bits into $m = \text{poly}(n)$ bits that “look random”. Such a function is called a *pseudorandom generator*. The question is how we can formalize the requirement that the output should “look random” in such a way that (a) the output can be used in place of the truly random bits in *any* **BPP** algorithm, and (b) such a generator exists.

Some candidate definitions for what it means for the random variable $X = G(U_d)$ to “look random” include the following:

- *Information-theoretic or statistical measures*: For example, we might measure entropy of $G(U_d)$, its statistical difference from the uniform distribution, or require pairwise independence. All of these fail one of the two criteria. For example, it is impossible for a

deterministic function to increase entropy from $O(\log n)$ to $\text{poly}(n)$. And it is easy to construct algorithms that fail when run using random bits that are only guaranteed to be pairwise independent.

- *Kolmogorov complexity*: a string x “looks random” if it is incompressible (cannot be generated by a Turing machine with a description of length less than $|x|$). An appealing aspect of this notion is that it makes sense of the randomness in a fixed string (rather than a distribution). Unfortunately, it is not suitable for our purposes. Specifically, if the function G is computable (which we certainly want) then all of its outputs have Kolmogorov complexity $d = O(\log n)$ (just hardwire the seed into the TM computing G), and hence are very compressible.
- *Computational indistinguishability*: this is the measure we will use. Intuitively, we say that a random variable X “looks random” if no *efficient* algorithm can distinguish X from a truly uniform random variable. Another perspective comes from the definition of statistical difference:

$$\Delta(X, Y) = \max_T |\Pr[X \in T] - \Pr[Y \in T]|.$$

With computational indistinguishability, we simply restrict the max to be taken only over “efficient” statistical tests T — that is, T ’s for which membership can be efficiently tested.

7.1.1 Computational Indistinguishability

Definition 7.1 (computational indistinguishability). Random variables X and Y taking values in $\{0, 1\}^m$ are (t, ε) *indistinguishable* if for every *nonuniform* algorithm T running in time at most t , we have

$$|\Pr[T(X) = 1] - \Pr[T(Y) = 1]| \leq \varepsilon$$

The left-hand side above is called also the *advantage* of T .

Recall that a nonuniform algorithm is an algorithm that may have some nonuniform advice hardwired in. (See Definition 3.10.) If the algorithm runs in time t we require that the advice string is of length at most t . Typically, to make sense of complexity measures like running time, it is necessary to use asymptotic notions, because a Turing machine can encode a huge lookup table for inputs of any bounded size in its transition function. However, for nonuniform algorithms, we can avoid doing so by using Boolean circuits as our nonuniform model of computation. Similarly to Fact 3.11, every nonuniform Turing machine algorithm running in time $t(n)$ can be simulated by a sequence of Boolean circuit C_n of size $\tilde{O}(t(n))$ and conversely every sequence of Boolean circuits of size $s(n)$ can be simulated by a nonuniform Turing machine running in time $\tilde{O}(s(n))$. Thus, to make our notation cleaner, from now on, by “nonuniform algorithm running in time t ,” we mean “Boolean circuit of size t ,” where we measure the size by the number of AND and OR gates in the circuit. (For convenience, we don’t count the inputs and negations in the circuit size.) Note also that in Definition 7.1 we have not specified whether the distinguisher is deterministic or randomized; this is because a probabilistic distinguisher achieving advantage greater than ε can be turned into a deterministic distinguisher achieving advantage greater than ε by nonuniformly

fixing the randomness. (This is another example of how “nonuniformity is more powerful than randomness,” like in Corollary 3.12.)

It is also of interest to study computational indistinguishability and pseudorandomness against uniform algorithms.

Definition 7.2 (uniform computational indistinguishability). Let X_m, Y_m be some sequences of random variables on $\{0, 1\}^m$ (or $\{0, 1\}^{\text{poly}(m)}$). For functions $t : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon : \mathbb{N} \rightarrow [0, 1]$, we say that $\{X_m\}$ and $\{Y_m\}$ are $(t(m), \varepsilon(m))$ *indistinguishable for uniform algorithms* if for all probabilistic algorithms T running in time $t(m)$, we have that

$$|\Pr[T(X_m) = 1] - \Pr[T(Y_m) = 1]| \leq \varepsilon(m)$$

for all sufficiently large m , where the probabilities are taken over X_m, Y_m and the random coin tosses of T .

We will focus on the nonuniform definition in this survey, but will mention results about the uniform definition as well.

7.1.2 Pseudorandom Generators

Definition 7.3. A deterministic function $G : \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (t, ε) *pseudorandom generator (PRG)* if

- (1) $d < m$, and
- (2) $G(U_d)$ and U_m are (t, ε) indistinguishable.

Also, note that we have formulated the definition with respect to nonuniform computational indistinguishability, but an analogous uniform definition can be given.

People attempted to construct pseudorandom generators long before this definition was formulated. Their generators were tested against a battery of statistical tests (e.g. the number of 1’s and 0’s are approximately the same, the longest run is of length $O(\log m)$, etc.), but these fixed set of tests provided no guarantee that the generators will perform well in an arbitrary application (e.g. in cryptography or derandomization). Indeed, most classical constructions (e.g. linear congruential generators, as implemented in the standard C library) are known to fail in some applications.

Intuitively, the above definition guarantees that the pseudorandom bits produced by the generator are as good as truly random bits for *all* efficient purposes (where efficient means time at most t). In particular, we can use such a generator for derandomizing any algorithm of running time at most t . For the derandomization to be efficient, we will also need the generator to be efficiently computable.

Definition 7.4. We say a sequence of generators $\{G_m : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m\}$ is *computable in time $t(m)$* if there is a *uniform and deterministic* algorithm M such that for every $m \in \mathbb{N}$ and $x \in \{0, 1\}^{d(m)}$, we have $M(m, x) = G_m(x)$ and $M(m, x)$ runs in time at most $t(m)$. In addition, $M(m)$ (with no second input) should output the value $d(m)$ in time at most $t(m)$.

Note that even when we define the pseudorandomness property of the generator with respect to nonuniform algorithms, the efficiency requirement refers to uniform algorithms. As usual, for readability, we will usually refer to a single generator $G : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$, with it being implicit that we are really discussing a family $\{G_m\}$.

Theorem 7.5. Suppose that for all m there exists an $(m, 1/8)$ pseudorandom generator $G : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$ computable in time $t(m)$. Then $\mathbf{BPP} \subset \bigcup_c \mathbf{DTIME}(2^{d(n^c)} \cdot (n^c + t(n^c)))$.

Proof. Let $A(x; r)$ be a **BPP** algorithm that on inputs x of length n , can be simulated by Boolean circuits of size at most n^c , using coin tosses r . Without loss of generality, we may assume that $|r| = n^c$. (It will often be notationally convenient to assume that the number of random bits used by an algorithm equals its running time or circuit size, so as to avoid an extra parameter. However, most interesting algorithms will only actually read and compute with a smaller number of these bits, so as to leave time available for computation. Thus, one should actually think of an algorithm as only reading a prefix of its random string r .)

The idea is to replace the random bits used by A with pseudorandom bits generated by G , use the pseudorandomness property to show that the algorithm will still be correct with high probability, and finally enumerate over all possible seeds to obtain a deterministic algorithm.

Claim 7.6. For every x of length n , $A(x; G(U_{d(n^c)}))$ errs with probability smaller than $1/2$.

Proof of claim: Suppose that there exists some x on which $A(x; G(U_{d(n^c)}))$ errs with probability at least $1/2$. Then $T(\cdot) = A(x, \cdot)$ is a Boolean circuit of size at most n^c that distinguishes $G(U_{d(n^c)})$ from U_{n^c} with advantage at least $1/2 - 1/3 > 1/8$. (Notice that we are using the input x as nonuniform advice; this is why we need the PRG to be pseudorandom against nonuniform tests.) \square

Now, enumerate over all seeds of length $d(n^c)$ and take a majority vote. There are $2^{d(n^c)}$ of them, and for each we have to run both G and A . \square

Notice that we can afford for the generator G to have running time $t(m) = \text{poly}(m)$ or even $t(m) = \text{poly}(m) \cdot 2^{O(d(m))}$ without affecting the time of the derandomization by than more than a polynomial amount. In particular, for this application, it is OK if the generator runs in more time than the tests it fools (which are time m in this theorem). That is, for derandomization, it suffices to have G that is mildly explicit according to the following definition:

Definition 7.7.

- (1) A generator $G : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$ is *mildly explicit* if it is computable in time $\text{poly}(m, 2^{d(m)})$.
 - (2) A generator $G : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$ is *fully explicit* if it is computable in time $\text{poly}(m)$.
-

These definitions are analogous to the notions of mildly explicit and fully explicit for expander graphs in Section 4.3. The truth table of a mildly explicit generator can be constructed in time polynomial in its size (which is of size $m \cdot 2^{d(m)}$), whereas a fully explicit generator can be evaluated in time polynomial in its input and output lengths (like the neighbor function of a fully explicit expander).

Theorem 7.5 provides a tradeoff between the seed length of the PRG and the efficiency of the derandomization. Let's look at some typical settings of parameters to see how we might simulate **BPP** in the different deterministic time classes (see Definition 3.1):

- (1) Suppose that for every constant $\varepsilon > 0$, there is an $(m, 1/8)$ mildly explicit PRG with seed length $d(m) = m^\varepsilon$. Then $\mathbf{BPP} \subset \bigcap_{\varepsilon > 0} \mathbf{DTIME}(2^{n^\varepsilon}) \stackrel{\text{def}}{=} \mathbf{SUBEXP}$. Since it is known that **SUBEXP** is a proper subset of **EXP**, this is already a nontrivial improvement on the current inclusion $\mathbf{BPP} \subset \mathbf{EXP}$ (Proposition 3.2).
- (2) Suppose that there is an $(m, 1/8)$ mildly explicit PRG with seed length $d(m) = \text{polylog}(m)$. Then $\mathbf{BPP} \subset \bigcup_c \mathbf{DTIME}(2^{\log^c m}) \stackrel{\text{def}}{=} \tilde{\mathbf{P}}$.
- (3) Suppose that there is an $(m, 1/8)$ mildly explicit PRG with seed length $d(m) = O(\log m)$. Then $\mathbf{BPP} = \mathbf{P}$.

Of course, all of these derandomizations are contingent on the question of whether PRGs exist. As usual, our first answer is yes but the proof is not very helpful—it is nonconstructive and thus does not provide for an efficiently computable PRG:

Proposition 7.8. For all $m \in \mathbb{N}$ and $\varepsilon > 0$, there exists a (non-explicit) (m, ε) pseudorandom generator $G : \{0, 1\}^d \rightarrow \{0, 1\}^m$ with seed length $d = O(\log m + \log(1/\varepsilon))$.

Proof. The proof is by the probabilistic method. Choose $G : \{0, 1\}^d \rightarrow \{0, 1\}^m$ at random. Now, fix a time m algorithm, T . The probability (over the choice of G) that T distinguishes $G(U_d)$ from U_m with advantage ε is at most $2^{-\Omega(2^d \varepsilon^2)}$, by a Chernoff bound. There are $2^{\text{poly}(m)}$ nonuniform algorithms running in time m (i.e. circuits of size m). Thus, union-bounding over all possible T , we get that the probability that there exists a T breaking G is at most $2^{\text{poly}(m)} 2^{-\Omega(2^d \varepsilon^2)}$, which is less than 1 for d being $O(\log m + \log(1/\varepsilon))$. \square

Note that putting together Proposition 7.8 and Theorem 7.5 gives us another way to prove that $\mathbf{BPP} \subset \mathbf{P/poly}$ (Corollary 3.12): just let the advice string be the truth table of an $(n^c, 1/8)$ PRG (which can be described by $2^{d(n^c)} \cdot n^c = \text{poly}(n)$ bits), and then use that PRG in the proof of Theorem 7.5 to derandomize the **BPP** algorithm. However, if you unfold both this proof and our previous proof (where we do error reduction and then fix the coin tosses), you will see that both proofs amount to essentially the same “construction”.

7.2 Cryptographic PRGs

The theory of computational pseudorandomness discussed in this chapter emerged from cryptography, where researchers sought a definition that would ensure that using pseudorandom bits instead of truly random bits (e.g. when encrypting a message) would retain security against all computationally feasible attacks. In this setting, the generator G is used by the honest parties and thus

should be very efficient to compute. On the other hand, the distinguisher T corresponds to an attack carried about by an adversary, and we want to protect against adversaries that may invest a lot of computational resources into trying to break the system. Thus, one is led to require that the pseudorandom generators be secure even against distinguishers with greater running time than the generator. The most common setting of parameters in the theoretical literature is that the generator should run in a fixed polynomial time, but the adversary can run in an arbitrary polynomial time.

Definition 7.9. A generator $G_m : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$ is a *cryptographic pseudorandom generator* if:

- (1) G_m is fully explicit. That is, there is a constant b such that G_m is computable in time m^b .
- (2) G_m is an $(m^{\omega(1)}, 1/m^{\omega(1)})$ PRG. That is, for every constant c , G_m is an $(m^c, 1/m^c)$ pseudorandom generator for all sufficiently large m .

Due to space constraints and the fact that such generators are covered in other texts (see the Chapter Notes and References), we will not do an in-depth study of cryptographic generators, but just survey what is known about them.

The first question to ask is whether such generators exist at all. It is not hard to show that cryptographic pseudorandom generators cannot exist unless $\mathbf{P} \neq \mathbf{NP}$, indeed unless $\mathbf{NP} \not\subseteq \mathbf{P/poly}$. (See Problem 7.3.) Thus, we do not expect to establish the existence of such generators unconditionally, and instead need to make some complexity assumption. While it would be wonderful to show that $\mathbf{NP} \not\subseteq \mathbf{P/poly}$ implies the existence of cryptographic pseudorandom generators, that too seems out of reach. However, we can base them on the very plausible assumption that there are functions that are easy to evaluate but hard to invert.

Definition 7.10. $f_n : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a *one-way function* if:

- (1) There is a constant b such that f_n is computable in time n^b for sufficiently large n .
- (2) For every constant c and every nonuniform algorithm A running in time n^c :

$$\Pr[A(f_n(U_n)) \in f_n^{-1}(f_n(U_n))] \leq \frac{1}{n^c}$$

for all sufficiently large n .

Assuming the existence of one-way functions seems stronger than the assumption $\mathbf{NP} \not\subseteq \mathbf{P/poly}$. For example, it is an *average-case* complexity assumption, as it requires that f is hard to invert when evaluated on random inputs. Nevertheless, there are a number of candidate functions believed to be one-way. The simplest is integer multiplication: $f_n(x, y) = x \cdot y$, where x and y are $n/2$ -bit numbers. Inverting this function amounts to the integer factorization problem, for which no efficient algorithm is known.

A classic and celebrated result in the foundations of cryptography is that cryptographic pseudorandom generators can be constructed from any one-way function:

Theorem 7.11. The following are equivalent:

- (1) One-way functions exist.
 - (2) There exist cryptographic pseudorandom generators with seed length $d(m) = m - 1$.
 - (3) For every constant $\varepsilon > 0$, there exist cryptographic pseudorandom generators with seed length $d(m) = m^\varepsilon$.
-

Corollary 7.12. If one-way functions exist, then $\mathbf{BPP} \subset \mathbf{SUBEXP}$.

What about getting a better derandomization? The proof of the above theorem is more general quantitatively. It takes any one-way function $f_\ell : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ and a parameter m , and constructs a generator $G_m : \{0, 1\}^{\text{poly}(\ell)} \rightarrow \{0, 1\}^m$. The fact that G_m is pseudorandom is proven by a reduction as follows. Assuming for contradiction that we have an algorithm T that runs in time t and distinguishes G_m from uniform with advantage ε , we construct an algorithm T' running in time $t' = t \cdot (m/\varepsilon)^{O(1)}$ inverting f_ℓ (say with probability $1/2$). If $t' \leq \text{poly}(\ell)$, then this contradicts the one-wayness of f , and hence we conclude that T cannot exist and G_m is a (t, ε) pseudorandom generator.

Quantitatively, if f_ℓ is hard to invert by algorithms running in time $s(\ell)$ and we take $m = 1/\varepsilon = s(\ell)^{o(1)}$, then we have $t' \leq s(\ell)$ for every $t = \text{poly}(m)$ and sufficiently large ℓ . Thus, viewing the seed length d of G_m as a function of m , we have $d(m) = \text{poly}(\ell) = \text{poly}(s^{-1}(m^{\omega(1)}))$, where $m^{\omega(1)}$ denotes any superpolynomial function of m .

Thus:

- If $s(\ell) = \ell^{\omega(1)}$, we can get seed length $d(m) = m^\varepsilon$ for any desired constant $\varepsilon > 0$ and $\mathbf{BPP} \subset \mathbf{SUBEXP}$ (as discussed above).
- If $s(\ell) = 2^{\ell^{\Omega(1)}}$ (as is plausible for the factoring one-way function), then we get seed length $d(m) = \text{poly}(\log m)$ and $\mathbf{BPP} \subset \tilde{\mathbf{P}}$.

But we cannot get seed length $d(m) = O(\log m)$, as needed for concluding $\mathbf{BPP} = \mathbf{P}$, from this result. Even for the maximum possible hardness $s(\ell) = 2^{\Omega(\ell)}$, we get $d(m) = \text{poly}(\log m)$. In fact, Problem 7.3 shows that it is impossible to have a cryptographic PRG with seed length $O(\log m)$ meeting Definition 7.9, where we require that G_m be pseudorandom against all $\text{poly}(m)$ -time algorithms. However, for derandomization we only need G_m to be pseudorandom against a fixed poly-time algorithm, e.g. running in time $t = m$, and we would get such generators with seed length $O(\log m)$ if the aforementioned construction could be improved to yield seed length $d = O(\ell)$ instead of $d = \text{poly}(\ell)$.

Open Problem 7.13. Given a one-way function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ that is hard to invert by algorithms running in time $s = s(\ell)$ and a constant c , is it possible to construct a fully explicit (t, ε) pseudorandom generator $G : \{0, 1\}^d \rightarrow \{0, 1\}^m$ with seed length $d = O(\ell)$ and pseudorandomness against time $t = s \cdot (\varepsilon/m)^{O(1)}$?

The best known seed length for such a generator is $d = \tilde{O}(\ell^3 \cdot \log(m/\varepsilon)/\log^2 s)$, which is $\tilde{O}(\ell^2)$ for the case that $s = 2^{\Omega(\ell)}$ and $m = 2^{\Omega(\ell)}$ as discussed above.

The above open problem has long been solved in the positive for one-way *permutations* $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$. In fact, the construction of pseudorandom generators from one-way permutations has a particularly simple description:

$$G_m(x, r) = (\langle x, r \rangle, \langle f(x), r \rangle, \langle f(f(x)), r \rangle, \dots, \langle f^{(m-1)}(x), r \rangle),$$

where $|r| = |x| = \ell$ and $\langle \cdot, \cdot \rangle$ denotes inner product modulo 2. One intuition for this construction is the following. Consider the sequence $(f^{(m-1)}(U_\ell), f^{(m-2)}(U_\ell), \dots, f(U_\ell), U_\ell)$. By the fact that f is hard to invert (but easy to evaluate) it can be argued that the $i + 1$ 'st component of this sequence is infeasible to predict from the first i components except with negligible probability. Thus, it is a computational analogue of a block source. The pseudorandom generator then is obtained by a computational analogue of block-source extraction, using the strong extractor $\text{Ext}(x, r) = \langle x, r \rangle$. The fact that the extraction works in this computational setting, however, is much more delicate and complex to prove than in the setting of extractors, and relies on a “local list-decoding algorithm” for the corresponding code (namely the Hadamard code). See Problems 7.12 and 7.13. (We will discuss local list decoding in Section 7.6.)

Pseudorandom Functions. It turns out that a cryptographic pseudorandom generator can be used to build an even more powerful object — a family of *pseudorandom functions*. This is a family of functions $\{f_s : \{0, 1\}^d \rightarrow \{0, 1\}\}_{s \in \{0, 1\}^d}$ such that (a) given the seed s , the function f_s can be evaluated in polynomial time, but (b) without the seed, it is infeasible to distinguish an oracle for f_s from an oracle to a truly random function. Thus in some sense, the d -bit truly random seed s is stretched to 2^d pseudorandom bits (namely the truth table of f_s)!

Pseudorandom functions have applications in several domains:

- *Cryptography:* When two parties share a seed s to a PRF, they effectively share a random function $f : \{0, 1\}^d \rightarrow \{0, 1\}$. (By definition, the function they share is indistinguishable from random by any poly-time 3rd party.) Thus, in order for one party to send a message m encrypted to the other, they can simply choose a random $r \xleftarrow{\text{R}} \{0, 1\}^d$, and send $(r, f_s(r) \oplus m)$. With knowledge of s , decryption is easy; simply calculate $f_s(r)$ and XOR it to the second part of the received message. However, the value $f_s(r) \oplus m$ would look essentially random to anyone without knowledge of s .

This is just one example; pseudorandom functions have vast applicability in cryptography.

- *Learning Theory:* Here, PRFs are used mainly to prove negative results. The basic paradigm in computational learning theory is that we are given a list of examples of a function’s behavior, $(x_1, f(x_1)), (x_2, f(x_2)), \dots, (x_k, f(x_k))$, where the x_i ’s are being selected randomly from some underlying distribution, and we would like to predict what the function’s value will be on a new data point x_{k+1} coming from the same distribution. Information-theoretically, correct prediction is possible after a small number of samples (with high probability), assuming that the function has a small description (e.g. is computable by a poly-sized circuit). However, it is computationally hard to predict the output of a PRF f_s on a new point x_{k+1} after seeing its value on k points (and this holds even if the algorithm gets to make “membership queries” — choose the evaluation points on its

own, in addition to getting random examples from some underlying distribution). Thus, PRFs provide examples of functions that are efficiently computable yet hard to learn.

- *Hardness of Proving Circuit Lower Bounds:* One main approach to proving $\mathbf{P} \neq \mathbf{NP}$ is to show that some $f \in \mathbf{NP}$ doesn't have polynomial size circuits (equivalently, $\mathbf{NP} \not\subseteq \mathbf{P/poly}$). This approach has had very limited success — the only superpolynomial lower bounds that have been achieved have been using very restricted classes of circuits (monotone circuits, constant depth circuits, etc). For general circuits, the best lower bound that has been achieved for a problem in \mathbf{NP} is $5n - o(n)$.

Pseudorandom functions have been used to help explain why existing lower-bound techniques have so far not yielded superpolynomial circuit lower bounds. Specifically, it has been shown that any sufficiently “constructive” proof of superpolynomial circuit lower bounds (one that would allow us to certify that a randomly chosen function has no small circuits) could be used to distinguish a pseudorandom function from truly random in subexponential time and thus invert any one-way function in subexponential time. This is known as the “Natural Proofs” barrier to circuit lower bounds.

7.3 Hybrid Arguments

In this section, we introduce a very useful proof method for working with computational indistinguishability, known as the *hybrid argument*. We use it to establish two important facts — that computational indistinguishability is preserved under taking multiple samples, and that pseudorandomness is equivalent to next-bit unpredictability.

7.3.1 Indistinguishability of Multiple Samples

The following proposition illustrates that computational indistinguishability behaves like statistical difference when taking many independent repetitions; the distance ε multiplies by at most the number of copies (cf. Lemma 6.3, Part 6). Proving it will introduce useful techniques for reasoning about computational indistinguishability, and will also illustrate how working with such computational notions can be more subtle than working with statistical notions.

Proposition 7.14. If random variables X and Y are (t, ε) indistinguishable, then for every $k \in \mathbb{N}$, X^k and Y^k are $(t, k\varepsilon)$ indistinguishable (where X^k represents k independent copies of X).

Note that when $t = \infty$, this follows from Lemma 6.3, Part 6; the challenge here is to show that the same holds even when we restrict to computationally bounded distinguishers.

Proof. We will prove the contrapositive: if there is an efficient algorithm T distinguishing X^k and Y^k with advantage greater than $k\varepsilon$, then there is an efficient algorithm T' distinguishing X and Y with advantage greater than ε . The difference in this proof from the corresponding result about statistical difference is that we need to preserve efficiency when going from T to T' . The algorithm T' will naturally use the algorithm T as a subroutine. Thus this is a *reduction* in the same spirit as reductions used elsewhere in complexity theory (e.g. in the theory of \mathbf{NP} -completeness).

Suppose that there exists a nonuniform time t algorithm T such that

$$\left| \Pr[T(X^k) = 1] - \Pr[T(Y^k) = 1] \right| > k\varepsilon \tag{7.1}$$

We can drop the absolute value in the above expression without loss of generality. (Otherwise we can replace T with its negation; recall that negations are free in our measure of circuit size.)

Now we will use a “hybrid argument.” Consider the hybrid distributions $H_i = X^{k-i}Y^i$, for $i = 0, \dots, k$. Note that $H_0 = X^k$ and $H_k = Y^k$. Then Inequality (7.1) is equivalent to

$$\sum_{i=1}^k \Pr[T(H_{i-1}) = 1] - \Pr[T(H_i) = 1] > k\varepsilon,$$

since the sum telescopes. Thus, there must exist some $i \in [k]$ such that $\Pr[T(H_{i-1}) = 1] - \Pr[T(H_i) = 1] > \varepsilon$, i.e.

$$\Pr[T(X^{k-i}XY^{i-1}) = 1] - \Pr[T(X^{k-i}YY^{i-1}) = 1] > \varepsilon.$$

By averaging, there exists some x_1, \dots, x_{k-i} and y_{k-i+2}, \dots, y_k such that

$$\Pr[T(x_1, \dots, x_{k-i}, X, y_{k-i+2}, \dots, y_k) = 1] - \Pr[T(x_1, \dots, x_{k-i}, Y, y_{k-i+2}, \dots, y_k) = 1] > \varepsilon.$$

Then, define $T'(z) = T(x_1, \dots, x_{k-i}, z, y_{k-i+2}, \dots, y_k)$. Note that T' is a nonuniform algorithm with advice $i, x_1, \dots, x_{k-i}, y_{k-i+2}, \dots, y_k$ hardwired in. Hardwiring these inputs costs nothing in terms of circuit size. Thus T' is a nonuniform time t algorithm such that

$$\Pr[T'(X) = 1] - \Pr[T'(Y) = 1] > \varepsilon,$$

contradicting the indistinguishability of X and Y . □

While the parameters in the above result behave nicely, with (t, ε) going to $(t, k\varepsilon)$, there are some implicit costs. First, the amount of nonuniform advice used by T' is larger than that used by T . This is hidden by the fact that we are using the same measure t (namely circuit size) to bound both the time and the advice length. Second, the result is meaningless for large values of k (e.g. $k = t$), because a time t algorithm cannot read more than t bits of the input distributions X^k and Y^k .

We note that there is an analogue of the above result for computational indistinguishability against *uniform* algorithms (Definition 7.2), but it is more delicate, because we cannot simply hardwire $i, x_1, \dots, x_{k-i}, y_{k-i+2}, \dots, y_k$ as advice. Indeed, a direct analogue of the proposition as stated is known to be false. We need to add the additional condition that the distributions X and Y are efficiently samplable. Then T' can choose $i \stackrel{R}{\leftarrow} [k]$ at random, and randomly sample $x_1, \dots, x_{k-i} \stackrel{R}{\leftarrow} X, y_{k-i+2}, \dots, y_k \stackrel{R}{\leftarrow} Y$.

7.3.2 Next-Bit Unpredictability

In analyzing the pseudorandom generators that we construct, it will be useful to work with a reformulation of the pseudorandomness property, which says that given a prefix of the output, it should be hard to predict the next bit much better than random guessing.

For notational convenience, we deviate from our usual conventions and write X_i to denote the i 'th bit of random variable X , rather than the i 'th random variable in some ensemble. We have:

Definition 7.15. Let X be a random variable distributed on $\{0, 1\}^m$. For $t \in \mathbb{N}$ and $\varepsilon \in [0, 1]$, we say that X is (t, ε) *next-bit unpredictable* if for every nonuniform probabilistic algorithm P running in time t and every $i \in [m]$, we have:

$$\Pr [P(X_1 X_2 \cdots X_{i-1}) = X_i] \leq \frac{1}{2} + \varepsilon,$$

where the probability is taken over X and the coin tosses of P .

Note that the uniform distribution $X \equiv U_m$ is $(t, 0)$ next-bit unpredictable for every t . Intuitively, if X is pseudorandom, it must be next-bit unpredictable, as this is just one specific kind of test one can perform on X . In fact the converse also holds, and that will be the direction we use.

Proposition 7.16. Let X be a random variable taking values in $\{0, 1\}^m$. If X is a (t, ε) pseudorandom, then X is $(t - O(1), \varepsilon)$ next-bit unpredictable. Conversely, if X is (t, ε) next-bit unpredictable, then it is $(t, m \cdot \varepsilon)$ pseudorandom.

Proof. Here U denotes an r.v. uniformly distributed on $\{0, 1\}^m$ and U_i denotes the i 'th bit of U .

pseudorandom \Rightarrow next-bit unpredictable. The proof is by reduction. Suppose for contradiction that X is not $(t - 3, \varepsilon)$ next-bit unpredictable, so we have a predictor $P : \{0, 1\}^{i-1} \rightarrow \{0, 1\}$ that succeeds with probability at least $1/2 + \varepsilon$. We construct an algorithm $T : \{0, 1\}^m \rightarrow \{0, 1\}$ that distinguishes X from U_m as follows:

$$T(x_1 x_2 \cdots x_m) = \begin{cases} 1 & \text{if } P(x_1 x_2 \cdots x_{i-1}) = x_i \\ 0 & \text{otherwise.} \end{cases}$$

T can be implemented with the same number of \wedge and \vee gates as P , plus 3 for testing equality (via the formula $(x \wedge y) \vee (\neg x \wedge \neg y)$).

next-bit unpredictable \Rightarrow pseudorandom. Also by reduction. Suppose X is not pseudorandom, so we have a nonuniform algorithm T running in time t s.t.

$$\Pr[T(X) = 1] - \Pr[T(U) = 1] > \varepsilon,$$

where we have dropped the absolute values without loss of generality as in the proof of Proposition 7.14.

We now use a hybrid argument. Define $H_i = X_1 X_2 \cdots X_i U_{i+1} U_{i+2} \cdots U_m$. Then $H_m = X$ and $H_0 = U$. We have:

$$\sum_{i=1}^m (\Pr[T(H_i) = 1] - \Pr[T(H_{i-1}) = 1]) > \varepsilon,$$

since the sum telescopes. Thus, there must exist an i such that

$$\Pr[T(H_i) = 1] - \Pr[T(H_{i-1}) = 1] > \varepsilon/m.$$

This says that T is more likely to output 1 when we put X_i in the i 'th bit than when we put a random bit U_i . We can view U_i as being X_i with probability $1/2$ and being \bar{X}_i with probability

$1/2$. The only advantage T has must be coming from the latter case, because in the former case, the two distributions are identical. Formally,

$$\begin{aligned}
& \varepsilon/m \\
& < \Pr[T(H_i) = 1] - \Pr[T(H_{i-1}) = 1] \\
& = \Pr[T(X_1 \cdots X_{i-1} X_i U_{i+1} \cdots U_m) = 1] \\
& \quad - \left(\frac{1}{2} \cdot \Pr[T(X_1 \cdots X_{i-1} X_i U_{i+1} \cdots U_m) = 1] + \frac{1}{2} \cdot \Pr[T(X_1 \cdots X_{i-1} \overline{X}_i U_{i+1} \cdots U_m) = 1] \right) \\
& = \frac{1}{2} \cdot (\Pr[T(X_1 \cdots X_{i-1} X_i U_{i+1} \cdots U_m) = 1] - \Pr[T(X_1 \cdots X_{i-1} \overline{X}_i U_{i+1} \cdots U_m) = 1]) .
\end{aligned}$$

This motivates the following next-bit predictor:

$P(x_1 x_2 \cdots x_{i-1})$:

- (1) Choose random bits $u_i, \dots, u_m \stackrel{R}{\leftarrow} \{0, 1\}$.
- (2) Compute $b = T(x_1 \cdots x_{i-1} u_i \cdots u_m)$.
- (3) If $b = 1$, output u_i , otherwise output \overline{u}_i .

The intuition is that T is more likely to output 1 when $u_i = x_i$ than when $u_i = \overline{x}_i$. Formally, we have:

$$\begin{aligned}
& \Pr[P(X_1 \cdots X_{i-1}) = X_i] \\
& = \frac{1}{2} \cdot (\Pr[T(X_1 \cdots X_{i-1} U_i U_{i+1} \cdots U_m) = 1 | U_i = X_i] + \Pr[T(X_1 \cdots X_{i-1} U_i U_{i+1} \cdots U_m) = 0 | U_i \neq X_i]) \\
& = \frac{1}{2} \cdot (\Pr[T(X_1 \cdots X_{i-1} X_i U_{i+1} \cdots U_m) = 1] + 1 - \Pr[T(X_1 \cdots X_{i-1} \overline{X}_i U_{i+1} \cdots U_m) = 1]) \\
& > \frac{1}{2} + \frac{\varepsilon}{m}.
\end{aligned}$$

Note that as described P runs in time $t + O(m)$. Recalling that we are using circuit size as our measure of nonuniform time, we can reduce the running time to t as follows. First, we may nonuniformly fix the coin tosses u_i, \dots, u_m of P while preserving its advantage. Then all P does is run T on $x_1 \cdots x_{i-1}$ concatenated with some fixed bits and either output what T does or its negation (depending on the fixed value of u_i). Fixing some input bits and negation can be done without increasing circuit size. Thus we contradict the next-bit unpredictability of X . \square

We note that an analogue of this result holds for uniform distinguishers and predictors, provided that we change the definition of next-bit predictor to involve a random choice of $i \stackrel{R}{\leftarrow} [m]$ instead of a fixed value of i , and change the time bounds in the conclusions to be $t - O(m)$ rather than $t - O(1)$ and t . (We can't do tricks like in the final paragraph of the proof.) In contrast to the multiple-sample indistinguishability result of Proposition 7.14, this result does not need X to be efficiently samplable for the uniform version.

7.4 Pseudorandom Generators from Average-Case Hardness

In Section 7.2, we surveyed cryptographic pseudorandom generators, which have numerous applications within and outside cryptography, including to derandomizing **BPP**. However, for deran-

domization, we can use generators with weaker properties. Specifically, Theorem 7.5 only requires $G : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$ such that:

- (1) G fools (nonuniform) distinguishers running in time m (as opposed to all $\text{poly}(m)$ -time distinguishers).
- (2) G is computable in time $\text{poly}(m, 2^{d(m)})$ (i.e. G is mildly explicit). In particular, *the PRG may take more time than the distinguishers it is trying to fool.*

Such a generator implies that every **BPP** algorithm can be derandomized in time $\text{poly}(n) \cdot 2^{d(\text{poly}(n))}$.

The benefit of studying such generators is that we can hope to construct them under weaker assumptions than used for cryptographic generators. In particular, a generator with the properties above no longer seems to imply $\mathbf{P} \neq \mathbf{NP}$, much less the existence of one-way functions. (The nondeterministic distinguisher that tests whether a string is an output of the generator by guessing a seed needs to evaluate the generator, which takes more time than the distinguishers are allowed.)

However, as shown in Problem 7.1, such generators still imply nonuniform circuit lower bounds for exponential time, something that is beyond the state of the art in complexity theory.

Our goal in the rest of this chapter is to construct generators as above from assumptions that are as weak as possible. In this section, we will construct them from boolean functions computable in exponential time that are hard on average (for nonuniform algorithms), and in the section after we will relax this to only require worst-case hardness.

7.4.1 Average-Case Hardness

A function is *hard on average* if it is hard to compute correctly on randomly chosen inputs. Formally:

Definition 7.17. For $s \in \mathbb{N}$ and $\delta \in [0, 1]$, we say that a Boolean function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is (s, δ) *average-case hard* if for all nonuniform probabilistic algorithms A running in time s ,

$$\Pr[A(X) = f(X)] \leq 1 - \delta,$$

where the probability is taken over X and the coin tosses of A .

Note that saying that f is (s, δ) hard for *some* $\delta > 0$ (possibly exponentially small) amounts to saying that f is *worst-case* hard.¹ Thus, we think of average-case hardness as corresponding to values of δ that are noticeably larger than zero, e.g. $\delta = 1/s^1$ or $\delta = 1/3$. Indeed, in this section we will take $\delta = 1/2 - \varepsilon$ for $\varepsilon = 1/s$. That is, no efficient algorithm can compute f much better than random guessing. A typical setting of parameters we use is $s = s(\ell)$ somewhere in range from $\ell^{\omega(1)}$ (slightly superpolynomial) to $s(\ell) = 2^{\alpha\ell}$ for a constant $\alpha > 0$. (Note that every function is computable by a nonuniform algorithm running in time roughly 2^ℓ , so we cannot take $s(\ell)$ to be any larger.) We will also require f to be computable in (uniform) time $2^{O(\ell)}$ so that our pseudorandom generator will be computable in time exponential in its seed length. The existence of such an average-case hard function may seem like a strong assumption, but in later sections we will see how to deduce it from a worst-case hardness assumption.

Now we show how to obtain a pseudorandom generator from average-case hardness.

¹For probabilistic algorithms, the “right” definition of worst-case hardness is actually that there exists an input x for which $\Pr[A(x) = f(x)] < 2/3$, where the probability is taken over the coin tosses of A . But for nonuniform algorithms two definitions can be shown to be roughly equivalent. See Definition 7.34 and the subsequent discussion.

Proposition 7.18. If $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is $(t, 1/2 - \varepsilon)$ average-case hard, then $G(x) = x \circ f(x)$ is a (t, ε) pseudorandom generator.

We omit the proof of this proposition, but it follows from Problem 7.5, Part 2 (by setting $m = 1$, $a = 0$, and $d = \ell$ in Theorem 7.24). Note that this generator includes its seed in its output. This is impossible for cryptographic pseudorandom generators, but is feasible (as shown above) when the generator can have more resources than the distinguishers it is trying to fool.

Of course, this generator is quite weak, stretching by only one bit. We would like to get many bits out. Here are two attempts:

- Use concatenation: Define $G(x_1 \cdots x_k) = x_1 \cdots x_k f(x_1) \cdots f(x_k)$. This is a $(t, k\varepsilon)$ pseudorandom generator because $G(U_{k\ell})$ consists of k independent samples of a pseudorandom distribution and thus computational indistinguishability is preserved by Proposition 7.14. Note that already here we are relying on *nonuniform* indistinguishability, because the distribution $(U_\ell, f(U_\ell))$ is not necessarily samplable (in time that is feasible for the distinguishers). Unfortunately, however, this construction does not improve the ratio between output length and seed length, which remains very close to 1.
- Use composition. For example, try to get two bits out using the same seed length by defining $G'(x) = G(G(x)_{1\dots\ell})G(x)_{\ell+1}$, where $G(x)_{1\dots\ell}$ denotes the first ℓ bits of $G(x)$. This works for cryptographic pseudorandom generators, but not for the generators we are considering here. Indeed, for the generator $G(x) = xf(x)$ of Proposition 7.18, we would get $G'(x) = xf(x)f(x)$, which is clearly not pseudorandom.

7.4.2 The Pseudorandom Generator

Our goal now is to show the following:

Theorem 7.19. For $s : \mathbb{N} \rightarrow \mathbb{N}$, suppose that there is a function $f \in \mathbf{E} = \mathbf{DTIME}(2^{O(\ell)})$ ² such that for every input length $\ell \in \mathbb{N}$, f is $(s(\ell), 1/2 - 1/s(\ell))$ average-case hard, where $s(\ell)$ is computable in time $2^{O(\ell)}$. Then for every $m \in \mathbb{N}$, there is a mildly explicit $(m, 1/m)$ pseudorandom generator $G : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$ with seed length $d(m) = O(s^{-1}(\text{poly}(m))^2 / \log m)$ that is computable in time $2^{O(d(m))}$.

Note that this is similar to the seed length $d(m) = \text{poly}(s^{-1}(\text{poly}(m)))$ mentioned in Section 7.2 for constructing cryptographic pseudorandom generators from one-way functions, but the average-case assumption is incomparable (and will be weakened further in the next section). In fact, it is known how to achieve a seed length $d(m) = O(s^{-1}(\text{poly}(m)))$, which matches what is known for constructing pseudorandom generators from one-way permutations as well as the converse implication of Problem 7.1. We will not cover that improvement here (see the Chapter Notes & References for pointers), but note that for the important case of hardness $s(\ell) = 2^{\Omega(\ell)}$, Theorem 7.19 achieves seed length $d(m) = O(O(\log m)^2 / \log m) = O(\log m)$ and thus $\mathbf{P} = \mathbf{BPP}$. More generally, we have:

² \mathbf{E} should be contrasted with the larger class $\mathbf{EXP} = \mathbf{DTIME}(2^{\text{poly}(\ell)})$. See Problem 7.2.

Corollary 7.20. Suppose that \mathbf{E} has a $(s(\ell), 1/2 - 1/s(\ell))$ average-case hard function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$.

- (1) If $s(\ell) = 2^{\Omega(\ell)}$, then $\mathbf{BPP} = \mathbf{P}$.
 - (2) If $s(\ell) = 2^{\ell^{\Omega(1)}}$, then $\mathbf{BPP} \subset \tilde{\mathbf{P}}$.
 - (3) If $s(\ell) = \ell^{\omega(1)}$, then $\mathbf{BPP} \subset \mathbf{SUBEXP}$.
-

The idea is to apply f repeatedly, but on *slightly dependent* inputs, namely ones that share very few bits. The sets of seed bits used for each output bit will be given by a *design*:

Definition 7.21. $S_1, \dots, S_m \subset [d]$ is an (ℓ, a) -*design* if

- (1) $\forall i, |S_i| = \ell$
 - (2) $\forall i \neq j, |S_i \cap S_j| \leq a$
-

We want lots of sets having small intersections over a small universe. We will use the designs established by Problem 3.2:

Lemma 7.22. For every constant $\gamma > 0$ and every $\ell, m \in \mathbb{N}$, there exists an (ℓ, a) -design $S_1, \dots, S_m \subset [d]$ with $d = O\left(\frac{\ell^2}{a}\right)$ and $a = \gamma \cdot \log m$. Such a design can be constructed deterministically in time $\text{poly}(m, d)$.

The important points are that intersection sizes are only logarithmic in the number of sets, and the universe size d is at most quadratic in the set size ℓ (and can be linear in ℓ in case we take $m = 2^{\Omega(\ell)}$).

Construction 7.23 (Nisan–Wigderson Generator). Given a function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ and an (ℓ, a) -design $S_1, \dots, S_m \subset [d]$, define the *Nisan–Wigderson generator* $G : \{0, 1\}^d \rightarrow \{0, 1\}^m$ as

$$G(x) = f(x|_{S_1})f(x|_{S_2}) \cdots f(x|_{S_m})$$

where if x is a string in $\{0, 1\}^d$ and $S \subset [d]$, then $x|_S$ is the string of length $|S|$ obtained from x by selecting the bits indexed by S .

Theorem 7.24. Let $G : \{0, 1\}^d \rightarrow \{0, 1\}^m$ be the Nisan–Wigderson generator based on a function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ and some (ℓ, a) design. If f is $(s, 1/2 - \varepsilon/m)$ average-case hard, then G is a (t, ε) pseudorandom generator, for $t = s - m \cdot a \cdot 2^a$.

Theorem 7.19 follows from Theorem 7.24 by setting $\varepsilon = 1/m$ and $a = \log m$, and observing that for $\ell = s^{-1}(m^3)$, then $t = s(\ell) - m \cdot a \cdot 2^a \geq m$, so we have an $(m, 1/m)$ pseudorandom generator. The seed length is $d = O(\ell^2 / \log m) = O(s^{-1}(\text{poly}(m))^2 / \log m)$.

Proof. Suppose G is not a (t, ε) pseudorandom generator. By Proposition 7.16, there is a nonuniform time t next-bit predictor P such that

$$\Pr[P(f(X|_{S_1})f(X|_{S_2}) \cdots f(X|_{S_{i-1}})) = f(X|_{S_i})] > \frac{1}{2} + \frac{\varepsilon}{m}, \quad (7.2)$$

for some $i \in [m]$. From P , we construct A that computes f with probability greater than $1/2 + \varepsilon/m$.

Let $Y = X|_{S_i}$. By averaging, we can fix all bits of $X|_{\overline{S_i}} = z$ (where $\overline{S_i}$ is the complement of S_i) such that the prediction probability remains greater than $1/2 + \varepsilon/m$ (over Y and the coin tosses of the predictor P). Define $f_j(y) = f(x|_{S_j})$ for $j \in \{1, \dots, i-1\}$. (That is, $f_j(y)$ forms x by placing y in the positions in S_i and z in the others, and then applies f to $x|_{S_j}$.) Then

$$\Pr_Y[P(f_1(Y) \cdots f_{i-1}(Y)) = f(Y)] > \frac{1}{2} + \frac{\varepsilon}{m}.$$

Note that $f_j(y)$ depends on only $|S_i \cap S_j| \leq a$ bits of y . Thus, we can compute each f_j with a look-up table, which we can include in the advice to our nonuniform algorithm. Indeed, every function on a bits can be computed by a boolean circuit of size at most $a \cdot 2^a$. (In fact, size at most $O(2^a/a)$ suffices.)

Then, defining $A(y) = P(f_1(y) \cdots f_{i-1}(y))$, we deduce that $A(y)$ can be computed with error probability smaller than $1/2 - \varepsilon/m$ in nonuniform time less than $t + m \cdot a \cdot 2^a = s$. This contradicts the hardness of f . Thus, we conclude G is an (m, ε) pseudorandom generator. \square

Some additional remarks on this proof:

- (1) This is a very general construction that works for any average-case hard function f . We only used $f \in \mathbf{E}$ to deduce G is computable in \mathbf{E} .
- (2) The reduction works for any nonuniform class of algorithms \mathcal{C} where functions of logarithmically many bits can be computed efficiently.

Indeed, in the next section we will use the same construction to obtain an *unconditional* pseudorandom generator fooling constant-depth circuits, and will later exploit the above “black-box” properties even further.

As mentioned earlier, the parameters of Theorem 7.24 have been improved in subsequent work, but the newer constructions do not have the clean structure of Nisan–Wigderson generator, where the seed of the generator is used to generate m random but correlated evaluation points, on which the average-case hard function f is evaluated. Indeed, each output bit of the improved generators depends on the entire truth-table of the function f , translating to a construction of significantly higher computational complexity. Thus the following remains an interesting open problem (which would have significance for hardness amplification as well as constructing pseudorandom generators):

Open Problem 7.25. For every $\ell, s \in \mathbb{N}$, construct an explicit generator $H : \{0, 1\}^{O(\ell)} \rightarrow \{0, 1\}^\ell$ with $m = s^{\Omega(1)}$ such that if f is $(s, 1/2 - 1/s)$ average-case hard and we define $G(x) = f(H_1(x))f(H_2(x)) \cdots f(H_m(x))$ where $H_i(x)$ denotes the i 'th component of $H(x)$, then G is an $(m, 1/m)$ pseudorandom generator.

7.4.3 Derandomizing Constant-depth circuits

Definition 7.26. An *unbounded fan-in circuit* $C(x_1, \dots, x_n)$ has input gates consisting of variables x_i , their negations $\neg x_i$, and the constants 0 and 1, as well as computation gates, which can compute the AND or OR of an unbounded number of other gates (rather than just 2, as in usual Boolean circuits).³ The *size* of such a circuit is the number of computation gates, and the *depth* is the maximum of length of a path from an input gate to the output gate.

\mathbf{AC}^0 is the class of functions $f : \{0, 1\}^* \rightarrow \{0, 1\}$ for which there exist constants c and k and a uniformly constructible sequence of unbounded fan-in circuits $(C_n)_{n \in \mathbb{N}}$ such that for all n , C_n has size at most n^c and depth at most k , and for all $x \in \{0, 1\}^n$, $C_n(x) = f(x)$. *Uniform constructibility* means that there is an efficient (e.g. polynomial-time) uniform algorithm M such that for all n , $M(1^n) = C_n$ (where 1^n denotes the number n in unary, i.e. a string of n 1's). \mathbf{BPAC}^0 defined analogously, except that C_n may have $\text{poly}(n)$ extra inputs, which are interpreted as random bits, and we require $\Pr_r[C_n(x, r) = f(x)] \geq 2/3$.

\mathbf{AC}^0 is one of the richest circuit classes for which we have superpolynomial lower bounds:

Theorem 7.27. For all constant $k \in \mathbb{N}$ and every $\ell \in \mathbb{N}$, the function $\text{PAR}_\ell : \{0, 1\}^\ell \rightarrow \{0, 1\}$ defined by $\text{PAR}_\ell(x_1, \dots, x_\ell) = \bigoplus_{i=1}^\ell x_i$ is $(s_k(\ell), 1/2 - 1/s_k(\ell))$ -average-case hard for nonuniform unbounded fan-in circuits of depth k and size $s_k(\ell) = 2^{\Omega(\ell^{1/k})}$.

The proof of this result is beyond the scope of this survey; see the Chapter Notes and References for pointers.

In addition to having an average-case hard function against \mathbf{AC}^0 , we also need that \mathbf{AC}^0 can compute arbitrary functions on a logarithmic number of bits.

Lemma 7.28. Every function $g : \{0, 1\}^a \rightarrow \{0, 1\}$ can be computed by a depth 2 circuit of size 2^a .

Using these two facts with the Nisan–Wigderson pseudorandom generator construction, we obtain the following pseudorandom generator for constant-depth circuits.

Theorem 7.29. For every constant k and every m , there exists a $\text{poly}(m)$ -time computable $(m, 1/m)$ -pseudorandom generator $G_m : \{0, 1\}^{\log^{O(k)} m} \rightarrow \{0, 1\}^m$ fooling unbounded fan-in circuits of depth k (and size m).

Proof. This is proven similarly to Theorems 7.19 and 7.24, except that we take $f = \text{PAR}_\ell$ rather than a hard function in \mathbf{E} , and we observe that the reduction can be implemented in a way that increases the depth by only an additive constant. Specifically, to obtain a pseudorandom generator fooling circuits of depth k and size m , we use the hardness of PAR_ℓ against unbounded fan-in circuits of depth $k' = k + 2$ and size m^2 , where $\ell = t_{k'}^{-1}(m^2) = O(\log^{k'} m)$. Then the seed length of G is $O(\ell^2/a) < O(\ell^2) = \log^{O(k)} m$.

We now follow the steps of the proof of Theorem 7.19 to go from an adversary T of depth k violating the pseudorandomness of G to a circuit A of depth k' calculating the parity function PAR_ℓ .

³Note that it is unnecessary to allow internal NOT gates, as these can always be pushed to the inputs via DeMorgan's Laws at no increase in size or depth.

If T has depth k , then it can be verified that the next-bit predictor P constructed in the proof of Proposition 7.16 also has depth k . (Recall that negations and constants can be propagated to the inputs so they do not contribute to the depth.) Next, in the proof of Theorem 7.24, we obtain A from P by $A(y) = P(f_1(y)f_2(y) \cdots f_{i-1}(y))$ for some $i \in \{1, \dots, m\}$ and where each f_i depends on at most a bits of y . Now we observe that A can be computed by a small constant-depth circuit (if P can). Specifically, applying Lemma 7.28 to each f_i , the size of A is at most $(m-1) \cdot 2^a + m = m^2$ and the depth of A is at most $k' = k + 2$. This contradicts the hardness of PAR_ℓ . \square

Corollary 7.30. $\text{BPAC}^0 \subset \tilde{\text{P}}$.

With more work, this can be strengthened to actually put BPAC^0 in $\widetilde{\text{AC}}^0$, i.e. uniform constant-depth circuits of quasipolynomial size. (The difficulty is that we use majority voting in the derandomization, but small constant-depth circuits cannot compute majority. However, they can compute an “approximate” majority, and this suffices.)

The above pseudorandom generator can also be used to give a quasipolynomial-time derandomization of the randomized algorithm we saw for approximately counting the number of satisfying assignments to a DNF formula (Theorem 2.34); see Problem 7.4.

Improving the running time of either of these derandomizations to polynomial is an intriguing open problem.

Open Problem 7.31. Show that $\text{BPAC}^0 = \text{AC}^0$ or even $\text{BPAC}^0 \subset \text{P}$.

Open Problem 7.32 (Open Problem 2.36, restated). Give a deterministic polynomial-time algorithm for approximately counting the number of satisfying assignments to a DNF formula.

We remark that it has recently been shown how to give an average-case AC^0 simulation of BPAC^0 (i.e. the derandomized algorithm is correct on *most* inputs); see Problem 7.5.

Another open problem is to construct similar, unconditional pseudorandom generators as Theorem 7.29 for circuit classes larger than AC^0 . A natural candidate is $\text{AC}^0[2]$, which is the same as AC^0 but augmented with unbounded-fan-in parity gates. There are known explicit functions $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ (e.g. MAJORITY) for which every $\text{AC}^0[2]$ circuit of depth k computing f has size at least $s_k(\ell) = 2^{\ell^{\Omega(1/k)}}$, but unfortunately the average-case hardness is much weaker than we need. These functions are only $(s_k(\ell), 1/2 - 1/o(\ell))$ -average-case hard, rather than $(s_k(\ell), 1/2 - 1/s_k(\ell))$ -average-case hard, so we can only obtain a small stretch using Theorem 7.24 and the following remains open.

Open Problem 7.33. For every constant k and every m , construct a (mildly) explicit $(m, 1/4)$ -pseudorandom generator $G_m : \{0, 1\}^{m^{o(1)}} \rightarrow \{0, 1\}^m$ fooling $\text{AC}^0[2]$ circuits of depth k and size m .

7.5 Worst-Case/Average-Case Reductions and Locally Decodable Codes

In the previous section, we saw how to construct pseudorandom generators from boolean functions that are very hard on average, where every nonuniform algorithm running in time t must err with probability greater than $1/2 - 1/t$ on a random input. Now we want to relax the assumption to refer to worst-case hardness, as captured by the following definition.

Definition 7.34. A function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is *worst-case hard for time t* if, for all nonuniform probabilistic algorithms A running in time t , there exists $x \in \{0, 1\}^\ell$ such that $\Pr[A(x) \neq f(x)] > 1/3$, where the probability is over the coin tosses of A .

Note that, for *deterministic* algorithms A , the definition simply says $\exists x A(x) \neq f(x)$. In the nonuniform case, restricting to deterministic algorithms is without loss of generality because we can always derandomize the algorithm using (additional) nonuniformity. Specifically, following the proof that $\mathbf{BPP} \subset \mathbf{P/poly}$, it can be shown that if f is worst-case hard for nonuniform deterministic algorithms running in time t , then it is worst-case hard for nonuniform probabilistic algorithms running in time t' for some $t' = \Omega(t/\ell)$.

A natural goal is to be able to construct an average-case hard function from a worst-case hard function. More formally, given a function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ that is worst-case hard for time $t = t(\ell)$, construct a function $\hat{f} : \{0, 1\}^{O(\ell)} \rightarrow \{0, 1\}$ such that \hat{f} is average-case hard for time $t' = t^{\Omega(1)}$. Moreover, we would like \hat{f} to be in \mathbf{E} if f is in \mathbf{E} . (Whether we can obtain a similar result for \mathbf{NP} is a major open problem, and indeed there are negative results ruling out natural approaches to doing so.)


Our approach to doing this will be via error-correcting codes. Specifically, we will show that if \hat{f} is the encoding of f in an appropriate kind of error-correcting code, then worst-case hardness of f implies average-case hardness of \hat{f} .

Specifically, we view f as a message of length $L = 2^\ell$, and apply an error-correcting code $\text{Enc} : \{0, 1\}^L \rightarrow \Sigma^{\hat{L}}$ to obtain $\hat{f} = \text{Enc}(f)$, which we view as a function $\hat{f} : \{0, 1\}^{\hat{\ell}} \rightarrow \Sigma$, where $\hat{\ell} = \log \hat{L}$. Pictorially:

$$\boxed{\text{message } f : \{0, 1\}^\ell \rightarrow \{0, 1\}} \longrightarrow \boxed{\text{Enc}} \longrightarrow \boxed{\text{codeword } \hat{f} : \{0, 1\}^{\hat{\ell}} \rightarrow \Sigma}.$$

(Ultimately, we would like $\Sigma = \{0, 1\}$, but along the way we will discuss larger alphabets.)

Now we argue the average-case hardness of \hat{f} as follows. Suppose, for contradiction, that \hat{f} is not δ average-case hard. By definition, there exists an efficient algorithm A with $\Pr[A(x) = \hat{f}(x)] > 1 - \delta$. We may assume that A is deterministic by fixing its coins. Then A may be viewed as a received word in $\Sigma^{\hat{L}}$, and our condition on A becomes $d_H(A, \hat{f}) < \delta$. So if Dec is a δ -decoding algorithm for Enc , then $\text{Dec}(A) = f$. By assumption A is efficient, so if Dec is efficient, then f may be efficiently computed everywhere. This would contradict our worst-case hardness assumption, assuming that $\text{Dec}(A)$ gives a time $t(\ell)$ algorithm for f . However, the standard notion of decoding requires reading all $2^{\hat{\ell}}$ values of the received word A and writing all 2^ℓ values of the message $\text{Dec}(A)$, and thus $\text{Time}(\text{Dec}(A)) \gg 2^\ell$. But every function on ℓ bits can be computed in nonuniform time 2^ℓ , and even in the uniform case we are mostly interested in $t(\ell) \ll 2^\ell$. To solve this problem we introduce the notion of local decoding.



Definition 7.35. A *local δ -decoding algorithm* for a code $\text{Enc} : \{0, 1\}^L \rightarrow \Sigma^{\hat{L}}$ is a probabilistic oracle algorithm Dec with the following property. Let $f : [L] \rightarrow \{0, 1\}$ be any message with associated codeword $\hat{f} = \text{Enc}(f)$, and let $g : [\hat{L}] \rightarrow \Sigma$ be such that $d_H(g, \hat{f}) < \delta$. Then for all $x \in [L]$ we have $\Pr[\text{Dec}^g(x) = f(x)] \geq 2/3$, where the probability is taken over the coins flips of Dec .

In other words, given oracle access to g , we want to efficiently compute any desired bit of f with high probability. So both the input (namely g) and the output (namely f) are treated implicitly; the decoding algorithm does not need to read/write either in its entirety. Pictorially:

This makes it possible to have sublinear-time (or even polylogarithmic-time) decoding. Also, we note that the bound of $2/3$ in the definition can be amplified in the usual way. Having formalized a notion of local decoding, we can now make our earlier intuition precise.

Proposition 7.36. Let Enc be an error-correcting code with local δ -decoding algorithm Dec that runs in nonuniform time at most t_{Dec} (meaning that Dec is a boolean circuit of size at most t_{Dec} equipped with oracle gates), and let f be worst-case hard for nonuniform time t . Then $\hat{f} = \text{Enc}(f)$ is (t', δ) average-case hard, where $t' = t/t_{\text{Dec}}$.

Proof. We do everything as explained before except with Dec^A in place of $\text{Dec}(A)$, and now the running time is at most $\text{Time}(\text{Dec}) \cdot \text{Time}(A)$. (We substitute each oracle gate in the circuit for Dec with the circuit for A .) \square

We note that the reduction in this proof does not use nonuniformity in an essential way. We used nonuniformity to fix the coin tosses of A , making it deterministic. To obtain a version for hardness against uniform probabilistic algorithms, the coin tosses of A can be chosen and fixed randomly instead. With high probability, the fixed coins will not increase A 's error by more than a constant factor (by Markov's Inequality); we can compensate for this by replacing the (t', δ) average-case hardness in the conclusion with, say, $(t', \delta/3)$ average-case hardness.

In light of the above proposition, our task is now to find an error-correcting code $\text{Enc} : \{0, 1\}^L \rightarrow \Sigma^{\hat{L}}$ with a local decoding algorithm. Specifically, we would like the following parameters.

- (1) We want $\hat{\ell} = O(\ell)$, or equivalently $\hat{L} = \text{poly}(L)$. This is because we measure hardness as a function of input length (which in turn translates to the relationship between output length and seed length of pseudorandom generators obtained via Theorem 7.19). In particular, when $t = 2^{\Omega(\ell)}$, we'd like to achieve $t' = 2^{\Omega(\hat{\ell})}$. Since $t' < t$ in Proposition 7.36, this is only possible if $\hat{\ell} = O(\ell)$.

- (2) We would like Enc to be computable in time $2^{O(\hat{\ell})} = \text{poly}(\hat{L})$, which is $\text{poly}(L)$ if we satisfy the requirement $\hat{L} = \text{poly}(L)$. This is because we want $f \in \mathbf{E}$ to imply $\hat{f} \in \mathbf{E}$.
- (3) We would like $\Sigma = \{0, 1\}$ so that \hat{f} is a boolean function, and $\delta = 1/2 - \varepsilon$ so that \hat{f} has sufficient average-case hardness for the pseudorandom generator construction of Theorem 7.24.
- (4) Since \hat{f} will be average-case hard against time $t' = t/t_{\text{Dec}}$, we would want the running time of Dec to be $t_{\text{Dec}} = \text{poly}(\ell, 1/\varepsilon)$ so that we can take $\varepsilon = t^{\Omega(1)}$ and still have $t' = t^{\Omega(1)}/\text{poly}(\ell)$.

Of course, achieving $\delta = 1/2 - \varepsilon$ is not possible with our current notion of local *unique* decoding (which is only harder than the standard notion of unique decoding), and thus in the next section we will focus on getting δ to be just a fixed constant. In Section 7.6, we will introduce a notion of local *list* decoding, which will enable decoding from distance $\delta = 1/2 - \varepsilon$.

In our constructions, it will be more natural to focus on the task of decoding *codeword* symbols rather than message symbols. That is, we replace the message f with the codeword \hat{f} in Definition 7.35 to obtain the following notion:

Definition 7.37 (locally correctible codes). ⁴ A *local δ -correcting algorithm* for a code $\mathcal{C} \subset \Sigma^{\hat{L}}$ is a probabilistic oracle algorithm Dec with the following property. Let $\hat{f} \in \mathcal{C}$ be any codeword, and let $g : [\hat{L}] \rightarrow \Sigma$ be such that $d_H(g, \hat{f}) < \delta$. Then for all $x \in [\hat{L}]$ we have $\Pr[\text{Dec}^g(x) = \hat{f}(x)] \geq 2/3$, where the probability is taken over the coin flips of Dec.

This implies the standard definition of locally decodable codes under the (mild) constraint that the message symbols are explicitly included in the codeword, as captured by the following definition (see also Problem 5.4).

Definition 7.38 (systematic encodings). An encoding algorithm $\text{Enc} : \{0, 1\}^L \rightarrow \mathcal{C}$ for a code $\mathcal{C} \subset \Sigma^{\hat{L}}$ is *systematic* if there is a polynomial-time computable function $I : [L] \rightarrow [\hat{L}]$ such that for all $f \in \{0, 1\}^L$, $\hat{f} = \text{Enc}(f)$, and all $x \in [L]$, we have $\hat{f}(I(x)) = f(x)$, where we interpret 0 and 1 as elements of Σ in some canonical way.

Informally, this means that the message f can be viewed as the restriction of the codeword \hat{f} to the coordinates in the image of I .

Lemma 7.39. If $\text{Enc} : \{0, 1\}^L \rightarrow \mathcal{C}$ is systematic and \mathcal{C} has a local δ -correcting algorithm running in time t , then Enc has a local δ -decoding algorithm (in the standard sense) running in time $t + \text{poly}(\log L)$.

Proof. If Dec_1 is the local corrector for \mathcal{C} and I the mapping in the definition of systematic encoding, then $\text{Dec}_2^g(x) = \text{Dec}_1^g(I(x))$ is a local decoder for Enc. □

⁴In the literature, these are often called *self-correctible codes*.

7.5.1 Local Decoding Algorithms

Hadamard Code. Recall the Hadamard code of message length m , which consists of the truth tables of all \mathbb{Z}_2 -linear functions $c : \{0, 1\}^m \rightarrow \{0, 1\}$ (Construction 5.12).

Proposition 7.40. The Hadamard code $\mathcal{C} \subset \{0, 1\}^{2^m}$ of message length m has a local $(1/4 - \varepsilon)$ -correcting algorithm running in time $\text{poly}(m, 1/\varepsilon)$.

Proof. We are given oracle access to $g : \{0, 1\}^m \rightarrow \{0, 1\}$ that is at distance less than $1/4 - \varepsilon$ from some (unknown) linear function c , and we want to compute $c(x)$ at an arbitrary point $x \in \{0, 1\}^m$. The idea is *random self-reducibility*: we can reduce computing c at an arbitrary point to computing c at uniformly random points, where g is likely to give the correct answer. Specifically, $c(x) = c(x \oplus r) \oplus c(r)$ for every r , and both $x \oplus r$ and r are uniformly distributed if we choose $r \stackrel{\text{R}}{\leftarrow} \{0, 1\}^m$. The probability that g differs from c at either of these points is less than $2 \cdot (1/4 - \varepsilon) = 1/2 - 2\varepsilon$. Thus $g(x \oplus r) \oplus g(r)$ gives the correct answer with probability noticeably larger than $1/2$. We can amplify this success probability by repetition. Specifically, we obtain the following local corrector:

Algorithm 7.41 (Local Corrector for Hadamard Code).

Input: an oracle $g : \{0, 1\}^m \rightarrow \{0, 1\}$, $x \in \{0, 1\}^m$, and a parameter $\varepsilon > 0$

- (1) Choose $r_1, \dots, r_t \stackrel{\text{R}}{\leftarrow} \{0, 1\}^m$, for $t = O(1/\varepsilon^2)$.
- (2) Query $g(r_i)$ and $g(r_i \oplus x)$ for each $i = 1, \dots, t$.
- (3) Output $\text{maj}_{1 \leq i \leq t} \{g(r_i) \oplus g(r_i \oplus x)\}$.

If $d_H(g, c) < 1/4 - \varepsilon$, then this algorithm will output $c(x)$ with probability at least $2/3$. □

This local decoding algorithm is optimal in terms of its decoding distance (arbitrarily close to $1/4$) and running time (logarithmic in the blocklength), but the problem is that the Hadamard code has exponentially small rate.

Reed–Muller Code. Recall that the q -ary Reed–Muller code of degree d and dimension m consists of all multivariate polynomials $p : \mathbb{F}_q^m \rightarrow \mathbb{F}_q$ of total degree at most d . (Construction 5.16.) This code has minimum distance $\delta = 1 - d/q$. Reed–Muller Codes are a common generalization of both Hadamard and Reed–Solomon codes, and thus we can hope that for an appropriate setting of parameters, we will be able to get the best of both kinds of codes. That is, we want to combine the efficient local decoding of the Hadamard code with the good rate of Reed–Solomon codes.

Theorem 7.42. The q -ary Reed–Muller Code of degree d and dimension m has a local $1/12$ -correcting algorithm running in time $\text{poly}(m, q)$ provided $d \leq q/9$ and $q \geq 36$.

Note the running time of the decoder is roughly the m 'th root of the block length $\hat{L} = q^m$. When $m = 1$, our decoder can query the entire string and we simply obtain a global decoding algorithm for Reed–Solomon Codes (which we already know how to achieve from Theorem 5.19).

But for large enough m , the decoder can only access a small fraction of the received word. (In fact, one can improve the running time to $\text{poly}(m, d, \log q)$, but the weaker result above is sufficient for our purposes.)

The key idea behind the decoder is to do restrictions to random *lines* in \mathbb{F}^m . The restriction of a Reed–Muller codeword to such a line is a Reed–Solomon codeword, and we can afford to run our global Reed–Solomon decoding algorithm on the line.

Formally, for $x, y \in \mathbb{F}^m$, we define the (*parameterized*) *line through x in direction y* as the function $\ell_{x,y} : \mathbb{F} \rightarrow \mathbb{F}^m$ given by $\ell_{x,y}(t) = x + ty$. Note that for every $a \in \mathbb{F}, b \in \mathbb{F} \setminus \{0\}$, the line $\ell_{x+ay, by}$ has the same set of points in its image as $\ell_{x,y}$; we refer to this set of points as an *unparameterized line*. When $y = 0$, the parameterized line contains only the single point x .

If $g : \mathbb{F}^m \rightarrow \mathbb{F}$ is any function and $\ell : \mathbb{F} \rightarrow \mathbb{F}^m$ is a line, then we use $g|_\ell$ to denote the *restriction of g to ℓ* , which is simply the composition $g \circ \ell : \mathbb{F} \rightarrow \mathbb{F}$. Note that if p is any polynomial of total degree at most d , then $p|_\ell$ is a (univariate) polynomial of degree at most d .

So we are given an oracle g of distance less than δ from some degree d polynomial $p : \mathbb{F}^m \rightarrow \mathbb{F}$, and we want to compute $p(x)$ for some $x \in \mathbb{F}^m$. We begin by choosing a random line ℓ through x . Every point of $\mathbb{F}^m \setminus \{x\}$ lies on exactly one parameterized line through x , so the points on ℓ (except x) are distributed uniformly at random over the whole domain, and thus g and p are likely to agree on these points. Thus we can hope to use the points on this line to reconstruct the value of $p(x)$. If δ is sufficiently small compared to the degree (e.g. $\delta = 1/3(d+1)$), we can simply interpolate the value of $p(x)$ from $d+1$ points on the line. This gives rise to the following algorithm.

Algorithm 7.43 (Local Corrector for Reed–Muller Code I).

Input: An oracle $g : \mathbb{F}^m \rightarrow \mathbb{F}$, an input $x \in \mathbb{F}^m$, and a degree parameter d

- (1) Choose $y \stackrel{\text{R}}{\leftarrow} \mathbb{F}^m$. Let $\ell = \ell_{x,y} : \mathbb{F} \rightarrow \mathbb{F}^m$ be the line through x in direction y .
 - (2) Query g to obtain $\beta_0 = g|_\ell(\alpha_0) = g(\ell(\alpha_0)), \dots, \beta_d = g|_\ell(\alpha_d) = g(\ell(\alpha_d))$ where $\alpha_0, \dots, \alpha_d \in \mathbb{F} \setminus \{0\}$ are any fixed points
 - (3) Interpolate to find the unique univariate polynomial q of degree at most d s.t. $\forall i, q(\alpha_i) = \beta_i$
 - (4) Output $q(0)$
-

Claim 7.44. If g has distance less than $\delta = 1/3(d+1)$ from some polynomial p of degree at most d , then Algorithm 7.43 will output $p(x)$ with probability greater than $2/3$.

Proof of claim: Observe that for all $x \in \mathbb{F}^m$ and $\alpha_i \in \mathbb{F} \setminus \{0\}$, $\ell_{x,y}(\alpha_i)$ is uniformly random in \mathbb{F}^m over the choice of $y \stackrel{\text{R}}{\leftarrow} \mathbb{F}^m$. This implies that for each i ,

$$\Pr_\ell [g|_\ell(\alpha_i) \neq p|_\ell(\alpha_i)] < \delta = \frac{1}{3(d+1)}.$$

By a union bound,

$$\Pr_\ell [\exists i, g|_\ell(\alpha_i) \neq p|_\ell(\alpha_i)] < (d+1) \cdot \delta = \frac{1}{3}.$$

Thus, with probability greater than $2/3$, we have $\forall i, q(\alpha_i) = p|_{\ell}(\alpha_i)$ and hence $q(0) = p(x)$. The running time of the algorithm is $\text{poly}(m, q)$. \square

We now show how to improve the decoder to handle a larger fraction of errors, up to distance $\delta = 1/12$. We alter Steps 2 and 3 in the above algorithm. In Step 2, instead of querying only $d + 1$ points, we query over *all* points in ℓ . In Step 3, instead of interpolation, we use a *global* decoding algorithm for Reed–Solomon codes to decode the univariate polynomial $p|_{\ell}$. Formally, the algorithm proceeds as follows.

Algorithm 7.45 (Local Corrector for Reed–Muller Codes II).

Input: An oracle $g: \mathbb{F}^m \rightarrow \mathbb{F}$, an input $x \in \mathbb{F}^m$, and a degree parameter d , where $q = |\mathbb{F}| \geq 36$ and $d \leq q/9$.

- (1) Choose $y \xleftarrow{R} \mathbb{F}^m$. Let $\ell = \ell_{x,y}: \mathbb{F} \rightarrow \mathbb{F}^m$ be the line through x in direction y .
 - (2) Query g at all points on ℓ to obtain $g|_{\ell}: \mathbb{F} \rightarrow \mathbb{F}$.
 - (3) Run the $1/3$ -decoder for the q -ary Reed–Solomon code of degree d on $g|_{\ell}$ to obtain the (unique) polynomial q at distance less than $1/3$ from $g|_{\ell}$ (if one exists).⁵
 - (4) Output $q(0)$.
-

Claim 7.46. If g has distance less than $\delta = 1/12$ from some polynomial p of degree at most d , and the parameters satisfy $q = |\mathbb{F}| \geq 36$, $d \leq q/9$, then Algorithm 7.45 will output $p(x)$ with probability greater than $2/3$.

Proof of claim: The expected distance (between $g|_{\ell}$ and $p|_{\ell}$) is small:

$$\mathbb{E}_{\ell}[d_H(g|_{\ell}, p|_{\ell})] < \frac{1}{q} + \delta = \frac{1}{36} + \frac{1}{12} = \frac{1}{9},$$

where the term $1/q$ is due to the fact that the point x is not random. Therefore, by Markov’s Inequality,

$$\Pr[d_H(g|_{\ell}, p|_{\ell}) \geq 1/3] \leq 1/3$$

Thus, with probability at least $2/3$, we have that $p|_{\ell}$ is the unique polynomial of degree at most d at distance less than $1/3$ from $g|_{\ell}$ and thus q must equal $p|_{\ell}$. \square

7.5.2 Low-Degree Extensions

Recall that to obtain locally decodable codes from locally correctible codes (as constructed above), we need to exhibit systematic encoding (Definition 7.38.) Thus, given $f: [L] \rightarrow \{0, 1\}$, we want to encode it as a Reed–Muller codeword $\hat{f}: [\hat{L}] \rightarrow \Sigma$ s.t.:

- The encoding time is $2^{O(\ell)} = \text{poly}(L)$

⁵ A $1/3$ -decoder for Reed–Solomon codes follows from the $(1 - 2\sqrt{d/q})$ list-decoding algorithm of Theorem 5.19. Since $1/3 \leq 1 - 2\sqrt{d/q}$, the list-decoder will produce a list containing all univariate polynomials at distance less than $1/3$, and since $1/3$ is smaller than half the minimum distance $(1 - d/q)$, there will be only one good decoding.

- $\hat{\ell} = O(\ell)$, or equivalently $\hat{L} = \text{poly}(L)$.
- The code is systematic in the sense of Definition 7.38.

Note that the usual encoding for Reed–Muller codes, where the message gives the coefficients of the polynomial, is not systematic. Instead the message should correspond to evaluations of the polynomial at certain points. Once we settle on the set of evaluation points, the task becomes one of interpolating the values at these points (given by the message) to a low-degree polynomial defined everywhere.

The simplest approach is to use the boolean hypercube as the set of evaluation points.

Lemma 7.47 (multilinear extension). For every $f: \{0, 1\}^\ell \rightarrow \{0, 1\}$ and every finite field \mathbb{F} , there exists a (unique) polynomial $\hat{f}: \mathbb{F}^\ell \rightarrow \mathbb{F}$ such that $\hat{f}|_{\{0,1\}^\ell} \equiv f$ and \hat{f} has degree at most 1 in each variable (and hence total degree at most ℓ).

Proof. We prove the existence of the polynomial \hat{f} . Define

$$\hat{f}(x_1, \dots, x_\ell) = \sum_{\alpha \in \{0,1\}^\ell} f(\alpha) \delta_\alpha(x)$$

for

$$\delta_\alpha(x) = \left(\prod_{i: \alpha_i=1} x_i \right) \left(\prod_{i: \alpha_i=0} (1 - x_i) \right)$$

Note that for $x \in \{0, 1\}^\ell$, $\delta_\alpha(x) = 1$ only when $\alpha = x$, therefore $\hat{f}|_{\{0,1\}^\ell} \equiv f$. We omit the proof of uniqueness. The bound on the individual degrees is by inspection. \square

Thinking of \hat{f} as an encoding of f , let's inspect the properties of this encoding.

- Since the total degree of the multilinear extension can be as large as ℓ , we need $q \geq 9\ell$ for the local corrector of Theorem 7.42 to apply.
- The encoding time is $2^{O(\ell)}$, as computing a single point of \hat{f} requires summing over 2^ℓ elements, and we have 2^ℓ points on which to compute \hat{f} .
- The code is systematic, since \hat{f} is an extension of f .
- However, the input length is $\hat{\ell} = \ell \log q = \Theta(\ell \log \ell)$, which is slightly larger than our target of $\hat{\ell} = O(\ell)$.

To solve the problem of the input length $\hat{\ell}$ in the multilinear encoding, we reduce the dimension of the polynomial \hat{f} by changing the embedding of the domain of f : Instead of interpreting $\{0, 1\}^\ell \subset \mathbb{F}^\ell$ as an embedding of the domain of f in \mathbb{F}^ℓ , we map $\{0, 1\}^\ell$ to \mathbb{H}^m for some subset $\mathbb{H} \subset \mathbb{F}$, and as such embed it in \mathbb{F}^m .

More precisely, we fix a subset $\mathbb{H} \subset \mathbb{F}$ of size $|\mathbb{H}| = \lceil \sqrt{q} \rceil$. Choose $m = \lceil \ell / \log |\mathbb{H}| \rceil$, and fix some efficient one-to-one mapping from $\{0, 1\}^\ell$ into \mathbb{H}^m . With this mapping, view f as a function $f: \mathbb{H}^m \rightarrow \mathbb{F}$.

Analogously to before, we have the following.

Lemma 7.48 (low-degree extension). For every finite field \mathbb{F} , $\mathbb{H} \subset \mathbb{F}$, $m \in \mathbb{N}$, and function $f : \mathbb{H}^m \rightarrow \mathbb{F}$, there exists a (unique) $\hat{f} : \mathbb{F}^m \rightarrow \mathbb{F}$ such that $\hat{f}|_{\mathbb{H}^m} \equiv f$ and \hat{f} has degree at most $|\mathbb{H}| - 1$ in each variable (and hence has total degree at most $m \cdot (|\mathbb{H}| - 1)$).

Using $|\mathbb{H}| = \lceil \sqrt{q} \rceil$, the total degree of \hat{f} is at most $d = \ell \sqrt{q}$. So we can apply the local corrector of Theorem 7.42, as long as $q \geq 81\ell^2$ (so that $d \leq q/9$). Inspecting the properties of \hat{f} as an encoding of f , we have:

- The input length is $\hat{\ell} = m \cdot \log q = \lceil \ell / \log |\mathbb{H}| \rceil \cdot \log q = O(\ell)$, as desired. (We can use a field of size 2^k for $k \in \mathbb{N}$, so that \hat{L} is a power of 2 and we incur no loss in encoding inputs to \hat{f} as bits.)
- The code is systematic as long as our mapping from $\{0, 1\}^{\hat{\ell}}$ to \mathbb{H}^m is efficient.

Note that not every polynomial of total degree at most $m \cdot (|\mathbb{H}| - 1)$ is the low-degree extension of a function $f : \mathbb{H}^m \rightarrow \mathbb{F}$, so the image of our encoding function $f \mapsto \hat{f}$ is only a *subcode* of the Reed–Muller code. This is not a problem, because any subcode of a locally correctible code is also locally correctible, and we can afford the loss in rate (all we need is $\hat{\ell} = O(\ell)$).

7.5.3 Putting It Together

Combining Theorem 7.42 with Lemmas 7.48, and 7.39, we obtain the following locally decodable code:

Proposition 7.49. For every $L \in \mathbb{N}$, there is an explicit code $\text{Enc} : \{0, 1\}^L \rightarrow \Sigma^{\hat{L}}$, with blocklength $\hat{L} = \text{poly}(L)$ and alphabet size $|\Sigma| = \text{poly}(\log L)$, that has a local $(1/12)$ -decoder running in time $\text{poly}(\log L)$.

Using Proposition 7.36, we obtain the following conversion from worst-case hardness to average-case hardness:

Proposition 7.50. If there exists $f : \{0, 1\}^{\ell} \rightarrow \{0, 1\}$ in \mathbf{E} that is worst-case hard against (non-uniform) time $t(\ell)$, then there exists $\hat{f} : \{0, 1\}^{O(\ell)} \rightarrow \{0, 1\}^{O(\log \ell)}$ in \mathbf{E} that is $(t'(\ell), 1/12)$ average-case hard for $t'(\ell) = t(\ell)/\text{poly}(\ell)$.

This differs from our original goal in two ways: \hat{f} is not Boolean, and we only get hardness $1/12$ (instead of $1/2 - \varepsilon$). The former concern can be remedied by concatenating the code of Proposition 7.49 with a Hadamard code, similarly to Problem 5.2. Note that the Hadamard code is applied on message space Σ , which is of size $\text{polylog}(L)$, so it can be $1/4$ -decoded by brute-force in time $\text{polylog}(L)$ (which is the amount of time already taken by our decoder).⁶ Using this, we obtain:

⁶Some readers may recognize this concatenation step as the same as applying the “Goldreich–Levin hardcore predicate” to \hat{f} . (See Problems 7.12 and 7.13.) However, for the parameters we are using, we do not need the power of these results, and can afford to perform brute-force unique decoding instead.

Theorem 7.51. For every $L \in \mathbb{N}$, there is an explicit code $\text{Enc} : \{0, 1\}^L \rightarrow \{0, 1\}^{\hat{L}}$ with blocklength $\hat{L} = \text{poly}(L)$ that has a local $(1/48)$ -decoder running in time $\text{poly}(\log L)$.

Theorem 7.52. If there exists $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ in \mathbf{E} that is worst-case hard against time $t(\ell)$, then there exists $\hat{f} : \{0, 1\}^{O(\ell)} \rightarrow \{0, 1\}$ in \mathbf{E} that is $(t'(\ell), 1/48)$ average-case hard, for $t'(\ell) = t(\ell)/\text{poly}(\ell)$.

An improved decoding distance can be obtained using Problem 7.7.

We note that the local decoder of Theorem 7.51 not only runs in time $\text{poly}(\log L)$, but also makes $\text{poly}(\log L)$ queries. For some applications (such as Private Information Retrieval, see Problem 7.6), it is important to have the number q of queries be as small as possible, ideally a constant. Using Reed–Muller codes of constant degree, it is possible to obtain constant-query locally decodable codes, but the blocklength will be $\hat{L} = \exp(L^{1/(q-1)})$. In a recent breakthrough, it was shown how to obtain constant-query locally decodable codes with blocklength $\hat{L} = \exp(L^{o(1)})$. Obtaining polynomial blocklength remains open.

Open Problem 7.53. Are there binary codes that are locally decodable with a constant number of queries (from constant distance $\delta > 0$) and blocklength polynomial in the message length?

7.5.4 Other Connections

As shown in Problem 7.6, locally decodable codes are closely related to protocols for *private information retrieval*. Another connection, and actually the setting in which these local decoding algorithms were first discovered, is to *program self-correctors*. Suppose we have a program for computing a function, such as the DETERMINANT, which happens to be a codeword in a locally decodable code (e.g. the determinant is a low-degree multivariate polynomial, and hence a Reed–Muller codeword). Then, even if this program has some bugs and gives the wrong answer on some small fraction of inputs, we can use the local decoding algorithm to obtain the correct answer on *all* inputs with high probability.

7.6 Local List Decoding and PRGs from Worst-Case Hardness

7.6.1 Hardness Amplification

In the previous section, we saw how to use locally decodable codes to convert worst-case hard functions into ones with constant average-case hardness (Theorem 7.52). Now our goal is to amplify this hardness (e.g. to $1/2 - \varepsilon$).

There are some generic techniques for hardness amplification. In these methods, we evaluate the function on many independent inputs. For example, consider f' that concatenates the evaluations of \hat{f} on k independent inputs:

$$f'(x_1, \dots, x_k) = (\hat{f}(x_1), \dots, \hat{f}(x_k)).$$

Intuitively, if \hat{f} is $1/12$ average-case hard, then f' should be $(1 - (11/12)^k)$ -average case hard because any efficient algorithm can solve each instance correctly with probability at most $11/12$. Proving this is nontrivial (because the algorithm trying to compute f' need not behave independently on the k instances), but there are Direct Product Theorems showing that the hardness does get amplified essentially as expected. Similarly, if we take the XOR on k independent inputs, the XOR Lemma says that the hardness approaches $1/2$ exponentially fast.

The main disadvantage of these approaches (for our purposes) is that the input length of f' is $k\ell$ while we aim for input length of $O(\ell)$. To overcome this problem, it is possible to use derandomized products, where we evaluate \hat{f} on *correlated* inputs instead of independent ones.

We will take a different approach, generalizing the notion and algorithms for locally decodable codes to locally *list*-decodable codes, and thereby directly construct \hat{f} that is $(1/2 - \varepsilon)$ -hard. Nevertheless, the study of hardness amplification is still of great interest, because it (or variants) can be employed in settings where doing a global encoding of the function is infeasible (e.g. for amplifying the average-case hardness of functions in complexity classes lower than \mathbf{E} , such as \mathbf{NP} , and for amplifying the security of cryptographic primitives). We remark that results on hardness amplification can be interpreted in a coding-theoretic language as well, as converting locally decodable codes with a small decoding distance into locally list-decodable codes with a large decoding distance. (See Section 8.2.3.)

7.6.2 Definition

We would like to formulate a notion of local *list*-decoding to enable us to have binary codes that are locally decodable from distances close to $1/2$. This is somewhat tricky to define — what does it mean to produce a “list” of decodings when only asked to decode a particular coordinate? Let g be our received word, and $\hat{f}_1, \hat{f}_2, \dots, \hat{f}_s$ the codewords that are close to g . One option would be for the decoding algorithm, on input x , to output a set of values $\text{Dec}^g(x) \subset \Sigma$ that is guaranteed to contain $\hat{f}_1(x), \hat{f}_2(x), \dots, \hat{f}_s(x)$ with high probability. However, this is not very useful; in the common case that $s \geq |\Sigma|$, the list could always be $\text{Dec}^g(x) = \Sigma$. Rather than outputting all of the values, we want to be able to specify to our decoder *which* $\hat{f}_i(x)$ to output. We do this with a two-phase decoding algorithm $(\text{Dec}_1, \text{Dec}_2)$, where both phases can be randomized.


- (1) Dec_1 , using g as an oracle and not given any other input other than the parameters defining the code, returns a list of advice strings a_1, a_2, \dots, a_s , which can be thought of as “labels” for each of the codewords close to g .
- (2) Dec_2 (again, using oracle access to g), takes input x and a_i , and outputs $\hat{f}_i(x)$.

The picture for Dec_2 is much like our old decoder, but it takes an extra input a_i corresponding to one of the outputs of Dec_1 :

More formally:

Definition 7.54. A *local δ -list-decoding algorithm* for a code Enc is a pair of probabilistic oracle algorithms $(\text{Dec}_1, \text{Dec}_2)$ such that for all received words g and all codewords $\hat{f} = \text{Enc}(f)$ with $d_H(\hat{f}, g) < \delta$, the following holds. With probability at least $2/3$ over $(a_1, \dots, a_s) \leftarrow \text{Dec}_1^g$, there exists an $i \in [s]$ such that

$$\forall x, \Pr[\text{Dec}_2^g(x, a_i) = f(x)] \geq 2/3.$$



Note that we don't explicitly require a bound on the list size s (to avoid introducing another parameter), but certainly it cannot be larger than the running time of Dec_1 .

As we did for locally (unique-)decodable codes, we can define a *local δ -list-correcting algorithm*, where Dec_2 should recover arbitrary symbols of the codeword \hat{f} rather than the message f . In this case, we don't require that for all j , $\text{Dec}_2^g(\cdot, a_j)$ is a codeword, or that it is close to g ; in other words, some of the a_j 's may be junk. Analogously to Lemma 7.39, a local δ -list-correcting algorithm implies local δ -list-decoding if the code is systematic.

Proposition 7.36 shows how locally decodable codes convert functions that are hard in the worst case to ones that are hard on average. The same is true for local list-decoding:

Proposition 7.55. Let Enc be an error-correcting code with local δ -list-decoding algorithm $(\text{Dec}_1, \text{Dec}_2)$ where Dec_2 runs in time at most t_{Dec} , and let f be worst-case hard for non-uniform time t . Then $\hat{f} = \text{Enc}(f)$ is (t', δ) average-case hard, where $t' = t/t_{\text{Dec}}$.

Proof. Suppose for contradiction that \hat{f} is not (t', δ) -hard. Then some nonuniform algorithm A running in time t' computes \hat{f} with error probability smaller than δ . But if Enc has a local δ list-decoding algorithm, then (with A playing the role of g) that means there exists a_i (one of the possible outputs of Dec_1^A), such that $\text{Dec}_2^A(\cdot, a_i)$ computes $f(\cdot)$ everywhere. Hardwiring a_i as advice, $\text{Dec}_2^A(\cdot, a_i)$ is a nonuniform algorithm running in time at most $\text{time}(A) \cdot \text{time}(\text{Dec}_2) \leq t$. \square

Note that, in contrast to Proposition 7.36, here we are using nonuniformity more crucially, in order to select the right function from the list of possible decodings. As we will discuss in Section 7.7.1, this use of nonuniformity is essential for “black-box” constructions, that do not exploit any structure in the hard function f or the adversary (A in the above proof). However, there are results on hardness amplification against uniform algorithms, which use structure in the hard function f (e.g. that it is complete for a complexity class like **E** or **NP**) to identify it among the list of decodings without any nonuniform advice.

7.6.3 Local List-Decoding Reed–Muller Codes

Theorem 7.56. There is a universal constant c such that the q -ary Reed–Muller code of degree d and dimension m over can be locally $(1 - \varepsilon)$ -list-corrected in time $\text{poly}(q, m)$ for $\varepsilon = c\sqrt{d/q}$.

Note that the distance at which list decoding can be done approaches 1 as $q/d \rightarrow \infty$. It matches the bound for list-decoding Reed–Solomon codes (Theorem 5.19) up to the constant c . Moreover, as the dimension m increases, the running time of the decoder ($\text{poly}(q, m)$) becomes much smaller than the block length ($q^m \cdot \log q$), at the price of a reduced rate ($\binom{m+d}{m}/q^m$).

Proof. Suppose we are given an oracle $g : \mathbb{F}^m \rightarrow \mathbb{F}$ that is $(1 - \varepsilon)$ close to some unknown polynomial $p : \mathbb{F}^m \rightarrow \mathbb{F}$, and that we are given an $x \in \mathbb{F}^m$. Our goal is describe two algorithms, Dec_1 and Dec_2 , where Dec_2 is able to compute $p(x)$ using a piece of Dec_1 's output (i.e. advice).

The advice that we will give to Dec_2 is the value of p on a single point. Dec_1 can easily generate a (reasonably small) list that contains one such point by choosing a random $y \in \mathbb{F}^m$, and outputting all pairs (y, z) , for $z \in \mathbb{F}$. More formally:

Algorithm 7.57 (Reed–Muller Local List-Decoder Dec_1).

Input: An oracle $g : \mathbb{F}^m \rightarrow \mathbb{F}$ and a degree parameter d

- (1) Choose $y \xleftarrow{\mathbb{R}} \mathbb{F}^m$
 - (2) Output $\{(y, z) : z \in \mathbb{F}\}$
-

This first-phase decoder is rather trivial in that it doesn't make use of the oracle access to the received word g . It is possible to improve both the running time and list size of Dec_1 by using oracle access to g , but we won't need those improvements below.

Now, the task of Dec_2 is to calculate $p(x)$, given the value of p on some point y . Dec_2 does this by looking at g restricted to the line through x and y , and using the list-decoding algorithm for Reed–Solomon Codes to find the univariate polynomials q_1, q_2, \dots, q_t that are close to g . If exactly one of these polynomials q_i agrees with p on the test point y , then we can be reasonably confident that $q_i(x) = p(x)$.

In more detail, the decoder works as follows:

Algorithm 7.58 (Reed–Muller Local List-Corrector Dec_2).

Input: An oracle $g : \mathbb{F}^m \rightarrow \mathbb{F}$, an input $x \in \mathbb{F}^m$, advice $(y, z) \in \mathbb{F}^m \times \mathbb{F}$, and a degree parameter d

- (1) Let $\ell = \ell_{x, y-x} : \mathbb{F} \rightarrow \mathbb{F}^m$ be the line through x and y (so that $\ell(0) = x$ and $\ell(1) = y$).
 - (2) Run the $(1 - \varepsilon/2)$ -list-decoder for Reed–Solomon Codes (Theorem 5.19) on $g|_{\ell}$ to get all univariate polys q_1, \dots, q_t that agree with $g|_{\ell}$ in greater than an $\varepsilon/2$ fraction of points.
 - (3) If there exists a unique i such that $q_i(1) = z$, output $q_i(0)$. Otherwise, fail.
-

Now that we have fully specified the algorithms, it remains to analyze them and show that they decode p correctly. Observe that it suffices to compute p on greater than an $11/12$ fraction of the points x , because then we can apply the unique local correcting algorithm of Theorem 7.42. Therefore, to finish the proof of the theorem we must prove the following.

Claim 7.59. Suppose that $g : \mathbb{F}^m \rightarrow \mathbb{F}$ has agreement greater than ε with a polynomial $p : \mathbb{F}^m \rightarrow \mathbb{F}$ of degree at most d . For at least half of the points $y \in \mathbb{F}^m$ the following holds for greater than an $11/12$ fraction of lines ℓ going through y :

- (1) $\text{agr}(g|_\ell, p|_\ell) > \varepsilon/2$.
 - (2) There does not exist any univariate polynomial q of degree at most d other than $p|_\ell$ such that $\text{agr}(g|_\ell, q) > \varepsilon/2$ and $q(y) = p(y)$.
-

Proof of claim: It suffices to show that Items 1 and 2 hold with probability 0.99 over the choice of a random point $y \stackrel{\mathbb{R}}{\leftarrow} \mathbb{F}^m$ and a random line ℓ through y ; then we can apply Markov's inequality to finish the job.

Item 1 holds by pairwise independence. If the line ℓ is chosen randomly, then the q points on ℓ are pairwise independent samples of \mathbb{F}^m . The expected agreement between $g|_\ell$ and $p|_\ell$ is simply the agreement between g and p , which is greater than ε by hypothesis. So by the Pairwise-Independent Tail Inequality (Prop. 3.28),

$$\Pr[\text{agr}(g|_\ell, p|_\ell) \leq \varepsilon/2] < \frac{1}{q \cdot (\varepsilon/2)^2},$$

which can be made smaller than .01 for a large enough choice of the constant c in $\varepsilon = c\sqrt{d/q}$.

To prove Item 2, we imagine first choosing the line ℓ uniformly at random from all lines in \mathbb{F}^m , and then choosing y uniformly at random from the points on ℓ (reparameterizing ℓ so that $\ell(1) = y$). Once we choose ℓ , we can let q_1, \dots, q_t be all polynomials of degree at most d , other than $p|_\ell$, that have agreement greater than $\varepsilon/2$ with $g|_\ell$. (Note that this list is independent of the parametrization of ℓ , i.e. if $\ell'(x) = \ell(ax + b)$ for $a \neq 0$ then $p|_{\ell'}$ and $q'_i(x) = q_i(ax + b)$ have agreement equal to $\text{agr}(p|_\ell, q_i)$.) By the list-decodability of Reed–Solomon Codes (Proposition 5.15), we have $t = O(\sqrt{q/d})$.

Now, since two distinct polynomials can agree in at most d points, when we choose a random point $y \stackrel{\mathbb{R}}{\leftarrow} \ell$, the probability that q_i and p agree at y is at most d/q . After reparameterization of ℓ so that $\ell(1) = y$, this gives

$$\Pr_y[\exists i : q_i(1) = p(1)] \leq t \cdot \frac{d}{q} = O\left(\sqrt{\frac{d}{q}}\right).$$

This can also be made smaller than 0.01 for large enough choice of the constant c (since we may assume $q/d > c^2$, else $\varepsilon \geq 1$ and the result holds vacuously). \square

\square

7.6.4 Putting it Together

To obtain a locally list-decodable (rather than list-correctible) code, we again use the low-degree extension (Lemma 7.48) to obtain a systematic encoding. As before, to encode messages of length

$\ell = \log L$, we apply Lemma 7.48 with $|\mathbb{H}| = \lceil \sqrt{q} \rceil$ and $m = \lceil \ell / \log |\mathbb{H}| \rceil$, for total degree $d \leq \sqrt{q} \cdot \ell$. To decode from a $1 - \varepsilon$ fraction of errors using Theorem 7.56, we need $c\sqrt{d/q} \leq \varepsilon$, which follows if $q \geq c^2 \ell^2 / \varepsilon^4$. This yields the following locally list-decodable codes:

Theorem 7.60. For every $L \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit code $\text{Enc} : \{0, 1\}^L \rightarrow \Sigma^{\hat{L}}$, with blocklength $\hat{L} = \text{poly}(L, 1/\varepsilon)$ and alphabet size $|\Sigma| = \text{poly}(\log L, 1/\varepsilon)$, that has a local $(1 - \varepsilon)$ -list-decoder running in time $\text{poly}(\log L, 1/\varepsilon)$.

Concatenating the code with a Hadamard code, similarly to Problem 5.2, we obtain:

Theorem 7.61. For every $L \in \mathbb{N}$ and $\varepsilon > 0$, there is an explicit code $\text{Enc} : \{0, 1\}^L \rightarrow \{0, 1\}^{\hat{L}}$ with blocklength $\hat{L} = \text{poly}(L, 1/\varepsilon)$ that has a local $(1/2 - \varepsilon)$ -list-decoder running in time $\text{poly}(\log L, 1/\varepsilon)$.

Using Proposition 7.55, we get the following hardness amplification result:

Theorem 7.62. For $s : \mathbb{N} \rightarrow \mathbb{N}$, suppose that there is a function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ in \mathbf{E} that is worst-case hard against nonuniform time $s(\ell)$, where $s(\ell)$ is computable in time $2^{O(\ell)}$, then there exists $\hat{f} : \{0, 1\}^{O(\ell)} \rightarrow \{0, 1\}$ in \mathbf{E} that is $(1/2 - 1/s'(\ell))$ average-case hard against (non-uniform) time $s'(\ell)$ for $s'(\ell) = t(\ell)^{\Omega(1)}/\text{poly}(\ell)$.

Combining this with Theorem 7.19 and Corollary 7.20, we get:

Theorem 7.63. For $s : \mathbb{N} \rightarrow \mathbb{N}$, suppose that there is a function $f \in \mathbf{E}$ such that for every input length $\ell \in \mathbb{N}$, f is worst-case hard for nonuniform time $s(\ell)$, where $s(\ell)$ is computable in time $2^{O(\ell)}$. Then for every $m \in \mathbb{N}$, there is a mildly explicit $(m, 1/m)$ pseudorandom generator $G : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$ with seed length $d(m) = O(s^{-1}(\text{poly}(m))^2 / \log m)$.

Corollary 7.64. For $s : \mathbb{N} \rightarrow \mathbb{N}$, suppose that there is a function $f \in \mathbf{E} = \mathbf{DTIME}(2^{O(\ell)})$ such that for every input length $\ell \in \mathbb{N}$, f is worst-case hard for nonuniform time $s(\ell)$. Then:

- (1) If $s(\ell) = 2^{\Omega(\ell)}$, then $\mathbf{BPP} = \mathbf{P}$.
 - (2) If $s(\ell) = 2^{\ell^{\Omega(1)}}$, then $\mathbf{BPP} \subset \tilde{\mathbf{P}}$.
 - (3) If $s(\ell) = \ell^{\omega(1)}$, then $\mathbf{BPP} \subset \mathbf{SUBEXP}$.
-

We note that the hypotheses in these results are simply asserting that there are problems in \mathbf{E} of high circuit complexity, which is quite plausible. Indeed, many common \mathbf{NP} -complete problems, such as SAT, are in \mathbf{E} and are commonly believed to have circuit complexity $2^{\Omega(\ell)}$ on inputs of length ℓ (though we seem very far from proving it). Thus, we have a “win-win” situation, either we can derandomize all of \mathbf{BPP} or SAT has significantly faster (nonuniform) algorithms than currently known.

Problem 7.1 establishes a converse to Theorem 7.63, showing that pseudorandom generators imply circuit lower bounds. The equivalence is fairly tight, except for the fact that Theorem 7.63 has seed length $d(m) = O(s^{-1}(\text{poly}(m))^2/\log m)$ instead of $d(m) = O(s^{-1}(\text{poly}(m)))$. It is known how to close this gap via a different construction, which is more algebraic and constructs PRGs directly from worst-case hard functions (see the Chapter Notes and References); a (positive) solution to Open Problem 7.25 would give a more modular and versatile construction. For Corollary 7.64, however, there is only a partial converse known. See Section 8.2.2.

Technical Comment. Consider Item 3 of Corollary 7.64, which assumes that there is a problem in \mathbf{E} of superpolynomial circuit complexity. This sounds similar to assuming that $\mathbf{E} \not\subseteq \mathbf{P}/\text{poly}$ (which is equivalent to $\mathbf{EXP} \not\subseteq \mathbf{P}/\text{poly}$, by Problem 7.2). However, the latter assumption is a bit weaker, because it only guarantees that there is a function $f \in \mathbf{E}$ and a function $s(\ell) = \ell^{\omega(1)}$ such that f has complexity at least $s(\ell)$ for *infinitely many* input lengths ℓ . Theorem 7.63 and Corollary 7.64 assume that f has complexity at least $s(\ell)$ for all ℓ ; equivalently f is not in $\mathbf{i.o.}\text{-P}/\text{poly}$, the class of functions that are computable by poly-sized circuits for infinitely many input lengths. We need the stronger assumptions because we want to build generators $G : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$ that are pseudorandom for all output lengths m , in order to get derandomizations of \mathbf{BPP} algorithms that are correct on all input lengths. However, there are alternate forms of these results, where the “infinitely often” is moved from the hypothesis to the conclusion. For example, if $\mathbf{E} \not\subseteq \mathbf{P}/\text{poly}$, we can conclude that $\mathbf{BPP} \subset \mathbf{i.o.}\text{-SUBEXP}$, where $\mathbf{i.o.}\text{-SUBEXP}$ denotes the class of languages having deterministic subexponential-time algorithms that are correct for infinitely many input lengths. Even though these “infinitely often” issues need to be treated with care for the sake of precision, it would be quite unexpected if the complexity of problems in \mathbf{E} and \mathbf{BPP} oscillated as a function of input length in such a strange way that they made a real difference.

7.7 Connections to Other Pseudorandom Objects

7.7.1 Black-Box Constructions

Similarly to our discussion after Theorem 7.19, the pseudorandom generator construction in the previous section is very general. The construction shows how to take *any* function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ and use it as a subroutine (oracle) to compute a generator $G^f : \{0, 1\}^d \rightarrow \{0, 1\}^m$ whose pseudorandomness can be related to the hardness of f . The only place that we use the fact that $f \in \mathbf{E}$ is to deduce that G^f is computable in \mathbf{E} . The reduction proving that G^f is pseudorandom is also very general. We showed how to take any T that distinguishes the output of $G^f(U_d)$ from U_m and use it as a subroutine (oracle) to build an efficient nonuniform algorithm Red such that Red^T computes f . The only place that we use the fact that T is itself an efficient nonuniform algorithm is to deduce that Red^T is an efficient nonuniform algorithm, contradicting the worst-case hardness of f .

Such constructions are called “black box,” because they treat the hard function f and the distinguisher T as black boxes (i.e. oracles), without using the code of the programs that compute f and T . As we will see, black-box constructions have significant additional implications. Thus, we formalize the notion of a black-box construction as follows:

Definition 7.65. Let $G^f : [D] \rightarrow [M]$ be a deterministic algorithm that is defined for every oracle $f : [L] \rightarrow \{0, 1\}$, let t, k be positive integers such that $k \leq t$, and let $\varepsilon > 0$. We say that G is a (t, k, ε) *black-box PRG construction* if there is a randomized oracle algorithm Red, running in time t , such that for every $f : [L] \rightarrow \{0, 1\}$ and $T : [M] \rightarrow \{0, 1\}$ such that if

$$\Pr[T(G^f(U_{[D]})) = 1] - \Pr[T(U_{[M]}) = 1] > \varepsilon,$$

then there is an advice string $z \in [K]$ such that

$$\forall x \in [L] \quad \Pr[\text{Red}^T(x, z) = f(x)] \geq 2/3,$$

where the probability is taken over the coin tosses of Red.

Note that we have separated the running time t of Red and the length k of its nonuniform advice into two separate parameters, and assume $k \leq t$ since an algorithm cannot read more than k bits in time t . When we think of Red as a nonuniform algorithm (like a boolean circuit), then we may as well think of these two parameters as being equal. (Recall that, up to $\text{polylog}(s)$ factors, being computable by a circuit of size s is equivalent to being computable by a uniform algorithm running in time s with s bits of nonuniform advice.) However, separating the two parameters is useful in order to isolate the role of nonuniformity, and to establish connections with the other pseudorandom objects we are studying.⁷

We note that if we apply a black-box pseudorandom generator construction with a function f that is actually hard to compute, then the result is indeed a pseudorandom generator:

Proposition 7.66. If G is a (t, k, ε) black-box PRG construction and f has nonuniform worst-case hardness at least s , then G^f is an $(s/\tilde{O}(t), \varepsilon)$ pseudorandom generator.

Now, we can rephrase the pseudorandom generator construction of Theorem 7.63 as follows:

Theorem 7.67. For every constant $\gamma > 0$, and every $\ell, m \in \mathbb{N}$, and every $\varepsilon > 0$, there is a (t, k, ε) black-box PRG construction $G^f : \{0, 1\}^d \rightarrow \{0, 1\}^m$ that is defined for every oracle $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$, with the following properties:

- (1) (Mild) explicitness: G^f is computable in uniform time $\text{poly}(m, 2^\ell)$ given an oracle for f .
 - (2) Seed length: $d = O((\ell + \log(1/\varepsilon))^2 / \log m)$.
 - (3) Reduction running time: $t = \text{poly}(m, 1/\varepsilon)$.
 - (4) Reduction advice length: $k = m^{1+\gamma} + O(\ell + \log(m/\varepsilon))$.
-

In addition to asserting the black-box nature of Theorem 7.63, the above is more general in that it allows ε to vary independently of m (rather than setting $\varepsilon = 1/m$), and gives a tighter bound on the length of the nonuniform advice than just $t = \text{poly}(m, 1/\varepsilon)$.

⁷Sometimes it is useful to allow the advice string z to also depend on the coin tosses of the reduction Red. By error reduction via $r = O(\ell)$ repetitions, such a reduction can be converted into one satisfying Definition 7.65 by sampling $r = O(\ell)$ sequences of coin tosses, but this blows up the advice length by a factor of r , which may be too expensive.

Proof Sketch. Given a function f , G^f encodes f in the locally list-decodable code of Theorem 7.61 (with decoding distance $1/2 - \varepsilon'$ for $\varepsilon' = \varepsilon/m$) to obtain $\hat{f} : \{0, 1\}^{\hat{\ell}} \rightarrow \{0, 1\}$ with $\hat{\ell} = O(\ell + \log(1/\varepsilon))$, and then computes the Nisan–Wigderson generator based on \hat{f} (Construction 7.23) and a $(\hat{\ell}, \gamma \log m)$ design. The seed length and mild explicitness follow from the explicitness and parameters of the design and code (Lemma 7.22 and Theorem 7.61). The running time and advice length of the reduction follow from inspecting the proofs of Theorems 7.61 and 7.19. Specifically, the running time of the Nisan–Wigderson reduction in the proof of Theorem 7.19 is $\text{poly}(m)$ (given the nonuniform advice) by inspection, and the running time of the local list-decoding algorithm is $\text{poly}(\ell, 1/\varepsilon') \leq \text{poly}(m, 1/\varepsilon)$. (We may assume that $m > \ell$, otherwise G^f need not have any stretch, and the conclusion is trivial.) The length of the advice from the locally list-decodable code consists of a pair $(y, z) \in \mathbb{F}^v \times \mathbb{F}$, where \mathbb{F} is a field of size $q = \text{poly}(\ell, 1/\varepsilon)$ and $v \log |\mathbb{F}| = \hat{\ell} = O(\ell + \log(1/\varepsilon))$. The Nisan–Wigderson reduction begins with the distinguisher-to-predictor reduction of Proposition 7.16, which uses $\log m$ bits of advice to specify the index i at which the predictor works and $m - i - 1$ bits for hardwiring the bits fed to the distinguisher in positions i, \dots, m . In addition, for $j = 1, \dots, i - 1$, the Nisan–Wigderson reduction nonuniformly hardwires a truth-table for the function $f_j(y)$ which depends on the at most $\gamma \cdot \log m$ bits of y selected by the intersection of the i 'th and j 'th sets in the design. These truth tables require at most $(i - 1) \cdot m^\gamma$ bits of advice. In total, the amount of advice used is at most

$$O(\ell + \log(1/\varepsilon)) + m - i - 1 + (i - 1) \cdot m^\gamma = m^{1+\gamma} + O(\ell + \log(1/\varepsilon)).$$

□

One advantage of a black-box construction is that it allows us to automatically “scale up” the pseudorandom generator construction. If we apply the construction to a function f that is not necessarily computable in \mathbf{E} , but in some higher complexity class, we get a pseudorandom generator G^f computable in an analogously higher complexity class. Similarly, if we want our pseudorandom generator to fool tests T computable by nonuniform algorithms in some higher complexity class, it suffices to use a function f that is hard against an analogously higher class.

For example, we get the following “nondeterministic” analogue of Theorem 7.63:

Theorem 7.68. For $s : \mathbb{N} \rightarrow \mathbb{N}$, suppose that there is a function $f \in \mathbf{NE} \cap \mathbf{co-NE}$ such that for every input length $\ell \in \mathbb{N}$, f is worst-case hard for nonuniform algorithms running in time $s(\ell)$ with an \mathbf{NP} oracle (equivalently, boolean circuits with SAT gates), where $s(\ell)$ is computable in time $2^{O(\ell)}$. Then for every $m \in \mathbb{N}$, there is a pseudorandom generator $G : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$ with seed length $d(m) = O(s^{-1}(\text{poly}(m))^2 / \log m)$ such that G is $(m, 1/m)$ -pseudorandom against nonuniform algorithms with an \mathbf{NP} oracle, and G is computable in nondeterministic time $2^{O(d(m))}$ (meaning that there is a nondeterministic algorithm that on input x , outputs $G(x)$ on at least one computation path and outputs either $G(x)$ or “fail” on all computation paths).

The significance of such generators is that they can be used for derandomizing \mathbf{AM} , which is a randomized analogue of \mathbf{NP} , defined as follows:

Definition 7.69. A language L is in **AM** iff there is a probabilistic polynomial-time verifier V and polynomials $m(n), p(n)$ such that for all inputs x of length n ,

$$\begin{aligned} x \in L &\Rightarrow \Pr_{r \leftarrow \{0,1\}^{m(n)}} [\exists y \in \{0,1\}^{p(n)} V(x, r, y) = 1] \geq 2/3, \\ x \notin L &\Rightarrow \Pr_{r \leftarrow \{0,1\}^{m(n)}} [\exists y \in \{0,1\}^{p(n)} V(x, r, y) = 1] \leq 1/3. \end{aligned}$$

Another (non-obviously!) equivalent definition of **AM** is the class of languages having constant-round interactive proof systems, where a computationally unbounded prover (“Merlin”) can convince a probabilistic polynomial-time verifier (“Arthur”) that an input x is in L through an interactive protocol with of $O(1)$ rounds of polynomial-length communication.

GRAPH NONISOMORPHISM is the most famous example of a language that is in **AM** but is not known to be in **NP**. Nevertheless, using Theorem 7.68 we can give evidence that GRAPH NONISOMORPHISM is in **NP**.

Corollary 7.70. If there is a function $f \in \mathbf{NE} \cap \mathbf{co-NE}$ that, on inputs of length ℓ , is worst-case hard for nonuniform algorithms running in time $2^{\Omega(\ell)}$ with an **NP** oracle, then **AM** = **NP**.

While the above complexity assumption may seem very strong, it is actually known to be weaker than the very natural assumption that exponential time $\mathbf{E} = \mathbf{DTIME}(2^{O(\ell)})$ is not contained in subexponential space $\cap_{\varepsilon > 0} \mathbf{DSpace}(2^{\varepsilon n})$.

As we saw in Section 7.4.3 on derandomizing constant-depth circuits, black-box constructions can also be “scaled down” to apply to lower complexity classes, provided that the construction G and/or reduction Red can be shown to be computable in a lower class (e.g. \mathbf{AC}^0).

7.7.2 Connections to Other Pseudorandom Objects

At first, it may seem that pseudorandom generators are of a different character than the other pseudorandom objects we have been studying. We require complexity assumptions to construct pseudorandom generators, and reason about them using the language of computational complexity (referring to efficient algorithms, reductions, etc.). The other objects we have been studying are all information-theoretic in nature, and our constructions of them have been unconditional.

The notion of black-box constructions will enable us to bridge this gap. Note that Theorem 7.67 is unconditional, and we will see that it, like all black-box constructions, has an information-theoretic interpretation. Indeed, we can fit black-box pseudorandom generator constructions into the list-decoding framework of Section 5.3 as follows:

Construction 7.71. Let $G^f : [D] \rightarrow [M]$ be an algorithm that is defined for every oracle $f : [n] \rightarrow \{0, 1\}$. Then, setting $N = 2^n$, define $\Gamma : [N] \times [D] \rightarrow [M]$, by

$$\Gamma(f, y) = G^f(y),$$

where we view the truth table of f as an element of $[N] \equiv \{0, 1\}^n$.

It turns out that if we allow the reduction unbounded running time (but still bound the advice length), then pseudorandom generator constructions have an exact characterization in our framework:

Proposition 7.72. Let G^f and Γ be as in Construction 7.71. Then G^f is an (∞, k, ε) black-box PRG construction iff for every $T \subset [M]$, we have

$$|\text{LIST}_\Gamma(T, \mu(T) + \varepsilon)| \leq K,$$

where $K = 2^k$.

Proof.

\Rightarrow . Suppose G^f is an (∞, k, ε) black-box PRG construction. Then f is in $\text{LIST}_\Gamma(T, \mu(T) + \varepsilon)$ iff T distinguishes $G^f(U_{[D]})$ from $U_{[M]}$ with advantage greater than ε . This implies that there exists a $z \in [K]$ such that $\text{Red}^T(\cdot, z)$ computes f everywhere. Thus, the number of functions f in $\text{LIST}_\Gamma(T, \mu(T) + \varepsilon)$ is bounded by the number of advice strings z , which is at most K .

\Leftarrow . Suppose that for every $T \subset [M]$, we have $L = |\text{LIST}_\Gamma(T, \mu(T) + \varepsilon)| \leq K$. Then we can define $\text{Red}^T(x, z) = f_z(x)$, where f_1, \dots, f_L are any fixed enumeration of the elements of $\text{LIST}_\Gamma(T, \mu(T) + \varepsilon)$. \square

Notice that this characterization of black-box PRG constructions (with reductions of unbounded running time) is the same as the one for averaging samplers (Proposition 5.30) and randomness extractors (Proposition 6.23). In particular, the black-box PRG construction of Theorem 7.67 is already a sampler and extractor of very good parameters:

Theorem 7.73. For every constant $\gamma > 0$, every $n \in \mathbb{N}$, $k \in [0, n]$, and every $\varepsilon > 0$, there is an explicit (k, ε) extractor $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with seed length $d = O(\log^2(n/\varepsilon)/\log k)$ and output length $m \geq k^{1-\gamma}$.

Proof Sketch. Without loss of generality, assume that n is a power of 2, namely $n = 2^\ell = L$. Let $G^f(y) : \{0, 1\}^d \rightarrow \{0, 1\}^m$ be the (t, k_0, ε_0) black-box PRG construction of Theorem 7.67 which takes a function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ and has $k_0 = m^{1+\gamma} + O(\ell + \log(m/\varepsilon_0))$, and let $\text{Ext}(f, y) = \Gamma(f, y) = G^f(y)$. By Propositions 7.72 and 6.23, Ext is a $(k_0 + \log(1/\varepsilon_0), 2\varepsilon_0)$ extractor. Setting $\varepsilon = 2\varepsilon_0$ and $k = k_0 + \log(1/\varepsilon_0)$, we have a (k, ε) extractor with output length

$$m = (k - O(\ell + \log(m/\varepsilon)))^{1-\gamma} \geq k^{1-\gamma} - O(\ell + \log(m/\varepsilon)).$$

We can increase the output length to $k^{1-\gamma}$ by increasing the seed length by $O(\ell + \log(m/\varepsilon))$. The total seed length then is

$$d = O\left(\frac{(\ell + \log(1/\varepsilon))^2}{\log k} + \ell + \log(m/\varepsilon)\right) = O\left(\frac{\log^2(n/\varepsilon)}{\log k}\right).$$

\square

The parameters of Theorem 7.73 are not quite as good as those of Theorem 6.36 and Corollary 6.39, as the output length is $k^{1-\gamma}$ rather than $(1 - \gamma)k$, and the seed length is only $O(\log n)$ when $k = n^{\Omega(1)}$. However, these settings of parameters are already sufficient for many purposes,

such as the simulation of **BPP** with weak random sources. Moreover, the extractor construction is much more direct than that of Theorem 6.36. Specifically, it is

$$\text{Ext}(f, y) = (\hat{f}(y|_{S_1}), \dots, \hat{f}(y|_{S_m})),$$

where \hat{f} is an encoding of f in a locally list-decodable code and S_1, \dots, S_m are a design. In fact, since Proposition 7.72 does not depend on the running time of the list-decoding algorithm, but only the amount of nonuniformity, we can use any $(1/2 - \varepsilon/2m, \text{poly}(m/\varepsilon))$ list-decodable code, which will only require an advice of length $O(\log(m/\varepsilon))$ to index into the list of decodings. In particular, we can use a Reed–Solomon code concatenated with a Hadamard code, as in Problem 5.2.

We now provide some additional intuition for why black-box pseudorandom generator constructions are also extractors. A black-box PRG construction G^f is designed to use a *computationally hard* function f (plus a random seed) to produce an output that is *computationally indistinguishable* from uniform. When we view it as an extractor $\text{Ext}(f, y) = G^f(y)$, we instead are feeding it a function f that is chosen randomly from a high min-entropy distribution (plus a random seed). This can be viewed as saying that f is *information-theoretically hard*, and from this stronger hypothesis, we are able to obtain the stronger conclusion that the output is *statistically indistinguishable* from uniform. The information-theoretic hardness of f can be formalized as follows: if f is sampled from a source F of min-entropy at least $k + \log(1/\varepsilon)$, then for every fixed function A (such as $A = \text{Red}^T$), the probability (over $f \leftarrow F$) that there exists a string z of length k such that $A(\cdot, z)$ computes f everywhere is at most ε . That is, a function generated with min-entropy larger than k is unlikely to have a description of length k (relative to any fixed “interpreter” A).

Similarly to black-box PRG constructions, we can also discuss converting worst-case hard functions to average-case hard functions in a black-box manner:

Definition 7.74. Let $\text{Amp}^f : [D] \rightarrow [q]$ be a deterministic algorithm that is defined for every oracle $f : [n] \rightarrow \{0, 1\}$. We say that Amp is a (t, k, ε) *black-box worst-case-to-average-case hardness amplifier* if there is a probabilistic oracle algorithm Red , called the *reduction*, running in time t such that for every function $g : [D] \rightarrow [q]$ such that

$$\Pr[g(U_{[D]}) = \text{Amp}^f(U_{[D]})] > 1/q + \varepsilon,$$

there is an advice string $z \in [K]$, where $K = 2^k$, such that

$$\forall x \in [n] \quad \Pr[\text{Red}^g(x, z) = f(x)] \geq 2/3,$$

where the probability is taken over the coin tosses of Red .

Note that this definition is almost identical to that of a locally $(1 - 1/q - \varepsilon)$ -list-decodable code (Definition 7.54), viewing Amp^f as $\text{Enc}(f)$, and Red^g as Dec_2^g . The only difference is that in the definition of locally list-decodable code, we require that there is a first-phase decoder Dec_1^g that efficiently produces a list of candidate advice strings (a property that is natural from a coding perspective, but is not needed when amplifying hardness against nonuniform algorithms). If we remove the constraint on the running time of Red , we simply obtain the notion of a $(1 - 1/q - \varepsilon, K)$ list-decodable code. By analogy, we can view black-box PRG constructions (with reductions of bounded running time) as simply being extractors (or averaging samplers) with a kind of efficient

local list-decoding algorithm (given by Red, again with an advice string that need not be easy to generate).

In addition to their positive uses illustrated above, black-box reductions and their information-theoretic interpretations are also useful for understanding the limitations of certain proof techniques. For example, we see that a black-box PRG construction $G^f : \{0, 1\}^d \rightarrow \{0, 1\}^m$ must have a reduction that uses $k \geq m - d - \log(1/\varepsilon) - 1$ bits of advice. Otherwise, by Propositions 7.72 and 6.23, we would obtain a $(k, 2\varepsilon)$ extractor that outputs m almost-uniform bits when given a source of min-entropy less than $k - d - 1$, which is impossible if $\varepsilon < 1/4$. Indeed, notions of black-box reduction have been used in other settings as well, most notably to produce a very fine understanding of the relations between different cryptographic primitives, meaning which ones can be constructed from each other via black-box constructions.

7.8 Exercises

Problem 7.1 (PRGs imply hard functions). Suppose that for every m , there exists a mildly explicit $(m, 1/m)$ pseudorandom generator $G_m : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$. Show that **E** has a function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ with nonuniform worst-case hardness $t(\ell) = \Omega(d^{-1}(\ell - 1))$. In particular, if $d(m) = O(\log m)$, then $t(\ell) = 2^{\Omega(\ell)}$ (Hint: look at a prefix of G 's output.)

Problem 7.2 (Equivalence of lower bounds for EXP and E). Show that **E** contains a function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ of circuit complexity $\ell^{\omega(1)}$ if and only if **EXP** does. (Hint: consider $f'(x_1 \cdots x_\ell) = f(x_1 \cdots x_{\ell^\varepsilon})$.)

Does the same argument work if we replace $\ell^{\omega(1)}$ with $2^{\ell^{\Omega(1)}}$? how about $2^{\Omega(\ell)}$?

Problem 7.3 (Limitations of Cryptographic Generators). (1) Prove that a cryptographic pseudorandom generator cannot have seed length $d(m) = O(\log m)$.
 (2) Prove that cryptographic pseudorandom generators (even with seed length $d(m) = m - 1$) imply **NP** $\not\subseteq$ **P/poly**.
 (3) Note where your proofs fail if we only require that G is an $(m^c, 1/m^c)$ pseudorandom generator for a fixed constant c .

Problem 7.4 (Deterministic Approximate Counting). Using the PRG for constant-depth circuits of Theorem 7.29, give deterministic quasipolynomial-time algorithms for the problems below. (The running time of your algorithms should be $2^{\text{poly}(\log n, \log(1/\varepsilon))}$, where n is the size of the circuit/formula given and ε is the accuracy parameter mentioned.)

- (1) Given a constant-depth circuit C and $\varepsilon > 0$, approximate the fraction of inputs x such that $C(x) = 1$ to within an additive error of ε .
- (2) Given a DNF formula φ and $\varepsilon > 0$, approximate the number of assignments x such that $\varphi(x) = 1$ to within a *multiplicative* fraction of $(1 + \varepsilon)$. You may restrict your attention to φ in which all clauses contain the same number of literals. (Hint: Study the randomized DNF counting algorithm of Theorem 2.34.)

Note that these are *not* decision problems, whereas classes such as **BPP** and **BPAC**₀ are classes of decision problems. One of the points of this problem is to show how derandomization can be used for other types of problems.

Problem 7.5 (Strong Pseudorandom Generators). By analogy with strong extractors, call a function $G : \{0, 1\}^d \rightarrow \{0, 1\}^m$ a (t, ε) *strong pseudorandom generator* iff the function $G'(x) = (x, G(x))$ is a (t, ε) pseudorandom generator.

- (1) Show that there do not exist strong cryptographic pseudorandom generators.
 - (2) Show that the Nisan–Wigderson generator (Theorem 7.24) is a strong pseudorandom generator.
 - (3) Suppose that for all constants $\alpha > 0$, there is a strong and fully explicit $(m, \varepsilon(m))$ pseudorandom generator $G : \{0, 1\}^{m^\alpha} \rightarrow \{0, 1\}^m$. Show that for every language $L \in \mathbf{BPP}$, there is a deterministic polynomial-time algorithm A such that for all n , $\Pr_{x \leftarrow \{0, 1\}^n} [A(x) \neq \chi_L(x)] \leq 1/2^n + \varepsilon(\text{poly}(n))$. That is, we get a *polynomial-time* average-case derandomization even though the seed length of G is $d(m) = m^\alpha$.
 - (4) Show that for every language $L \in \mathbf{BPAC}^0$, there is an **AC**⁰ algorithm A such that $\Pr_{x \leftarrow \{0, 1\}^n} [A(x) \neq \chi_L(x)] \leq 1/n$. (Warning: be careful about error reduction.)
-

Problem 7.6 (Private Information Retrieval). The goal of *private information retrieval* is for a user to be able to retrieve an entry of a remote database in such a way that the server holding the database *learns nothing* about which database entry was requested. A trivial solution is for the server to send the user the entire database, in which case the user does not need to reveal anything about the entry desired. We are interested in solutions that involve much less communication. One way to achieve this is through replication.⁸ Formally, in a q -server *private information-retrieval (PIR) scheme*, an arbitrary database $D \in \{0, 1\}^n$ is duplicated at q non-communicating servers. On input an index $i \in [n]$, the *user algorithm* U tosses some coins r and outputs queries $(x_1, \dots, x_q) = U(i, r)$, and sends x_j to the j 'th server. The j 'th server algorithm S_j returns an answer $y_j = S_j(x_j, D)$. The user then computes its output $U(i, r, x_1, \dots, x_q)$, which should equal D_i , the i 'th bit of the database. For privacy, we require that the distribution of each query x_j (over the choice of the random coin tosses r) is the same regardless of the index i being queried.

It turns out that q -query locally decodable codes and q -server PIR are essentially equivalent. This equivalence is proven using the notion of *smooth codes*. A code $\text{Enc} : \{0, 1\}^n \rightarrow \Sigma^{\hat{n}}$ is a q -query *smooth code* if there is a probabilistic oracle algorithm Dec such that for every message x and every $i \in [n]$, we have $\Pr[\text{Dec}^{\text{Enc}(x)}(i) = x_i] = 1$ and Dec makes q nonadaptive queries to its oracle, each of which is uniformly distributed in $[\hat{n}]$. Note that the oracle in this definition is a valid codeword, with no corruptions. Below you will show that smooth codes imply locally decodable codes and PIR schemes; converses are also known (after making some slight relaxations to the definitions).

⁸Another way is through computational security, where we only require that it be *computationally infeasible* for the database to learn something about the entry requested.

- (1) Show that the decoder for a q -query smooth code is also a local $(1/3q)$ -decoder for Enc.
- (2) Show that every q -query smooth code $\text{Enc} : \{0, 1\}^n \rightarrow \Sigma^{\hat{n}}$ gives rise to a q -server PIR scheme in which the user and servers communicate at most $q \cdot (\log \hat{n} + \log |\Sigma|)$ bits for each database entry requested.
- (3) Using the Reed-Muller code, show that there is a $\text{polylog}(n)$ -server PIR scheme with communication complexity $\text{polylog}(n)$ for n -bit databases. That is, the user and servers communicate at most $\text{polylog}(n)$ bits for each database entry requested. (For constant q , the Reed-Muller code with an optimal systematic encoding as in Problem 5.4 yields a q -server PIR with communication complexity $O(n^{1/(q-1)})$.)

Problem 7.7 (Better Local Decoding of Reed–Muller Codes). Show that for every constant $\varepsilon > 0$, there is a constant $\gamma > 0$ such that there is a local $(1/2 - \varepsilon)$ -decoding algorithm for the q -ary Reed-Muller code of degree d and dimension m , provided that $d \leq \gamma q$. (Here we are referring to unique decoding, not list decoding.) The running time of the decoder should be $\text{poly}(m, q)$.

Problem 7.8 (Hitting-Set Generators). A set $H_m \subset \{0, 1\}^m$ is a (t, ε) *hitting set* if for every nonuniform algorithm T running in time t that accepts greater than an ε fraction of m -bit strings, T accepts at least one element of H_m .

- (1) Show that if, for every m , we can construct an $(m, 1/2)$ hitting set H_m in time $s(m) \geq m$, then $\mathbf{RP} \subset \bigcup_c \mathbf{DTIME}(s(n^c))$. In particular, if $s(m) = \text{poly}(m)$, then $\mathbf{RP} = \mathbf{P}$.
- (2) Show that if there is a (t, ε) pseudorandom generator $G_m : \{0, 1\}^d \rightarrow \{0, 1\}^m$ computable in time s , then there is a (t, ε) hitting set H_m constructible in time $2^d \cdot s$.
- (3) Show that if, for every m , we can construct an $(m, 1/2)$ hitting set H_m in time $s(m) = \text{poly}(m)$, then $\mathbf{BPP} = \mathbf{P}$. (Hint: this can be proven in two ways. One uses Problem 3.1 and the other uses a variant of Problem 7.1 together with Corollary 7.64. How do the parameters for general $s(m)$ compare?)
- (4) Define the notion of a (t, k, ε) black-box construction of hitting set-generators, and show that, when $t = \infty$, such constructions are equivalent to constructions of *dispersers* (Definition 6.19).

Problem 7.9 (PRGs vs. Uniform Algorithms \Rightarrow Average-Case Derandomization). For functions $t : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon : \mathbb{N} \rightarrow [0, 1]$, we say that a sequence $\{G_m : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m\}$ of is a $(t(m), \varepsilon(m))$ pseudorandom generator against uniform algorithms iff the ensembles $\{G(U_{d(m)})\}_{m \in \mathbb{N}}$ and $\{U_m\}_{m \in \mathbb{N}}$ are uniformly computationally indistinguishable (Definition 7.2).

Suppose that we have a mildly explicit $(m, 1/m)$ pseudorandom generator against uniform algorithms that has seed length $d(m)$. Show that for every language L in \mathbf{BPP} , there exists a deterministic algorithm A running in time $2^{d(\text{poly}(n))} \cdot \text{poly}(n)$ on inputs of length n such that:

- (1) $\Pr [A(X_n) = L(X_n)] \geq 1 - 1/n^2$, where $X_n \stackrel{\text{R}}{\leftarrow} \{0, 1\}^n$ and $L(\cdot)$ is the characteristic function of L . (The exponent of 2 in n^2 is arbitrary, and can be replaced by any constant.) Hint: coming up with the algorithm A is the “easy” part; proving that it works well is a bit trickier.
- (2) $\Pr [A(X_n) = L(X_n)] \geq 1 - 1/n^2$, for any random variable X_n distributed on $\{0, 1\}^n$ that is samplable in time n^2 .

Problem 7.10 (PRGs are Necessary for Derandomization). (1) Call a function $G : \{0, 1\}^d \rightarrow \{0, 1\}^m$ a (t, ℓ, ε) pseudorandom generator against bounded-nonuniformity algorithms iff for every probabilistic algorithm T that has a program of length at most ℓ and that runs in time at most t on inputs of length n , we have

$$|\Pr[T(G(U_d)) = 1] - \Pr[T(U_m) = 1]| \leq \varepsilon.$$

Consider the promise problem Π whose YES instances are truth tables of functions $G : \{0, 1\}^d \rightarrow \{0, 1\}^m$ that are $(m, \log m, 1/m)$ pseudorandom generators against bounded-nonuniformity algorithms, and whose NO instances are truth tables of functions that are not $(m, \log m, 2/m)$ pseudorandom generators against bounded-nonuniformity algorithms. (Here m and d are parameters determined by the input instance G .) Show that Π is in **prBPP**.

- (2) Using Problem 2.11, show that if **prBPP** = **prP**, then there is a mildly explicit $(m, 1/m)$ pseudorandom generator against uniform algorithms with seed length $O(\log m)$. (See Problem 7.9 for the definition. It turns out that the hypothesis **prBPP** = **prP** here can be weakened to obtain an equivalence between PRGs vs. uniform algorithms and average-case derandomization of **BPP**.)

Problem 7.11 (Composition). For simplicity in this problem, only consider constant t in this problem (although the results do have generalizations to growing $t = t(\ell)$).

- (1) Show that if $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ is a one-way permutation, then for any constant t , $f^{(t)}$ is a one-way permutation, where $f^{(t)}(x) \stackrel{\text{def}}{=} \underbrace{f(f(\dots f(x)))}_t$.
- (2) Show that the above fails for one-way functions. That is, assuming that there exists a one-way function g , construct a one-way function f which doesn't remain one way under composition. (Hint: for $|x| = |y| = \ell/2$, set $f(x, y) = 1^{|g(y)|}g(y)$ unless $x \in \{0^\ell, 1^\ell\}$.)
- (3) Show that if G is a cryptographic pseudorandom generator with seed length $d(m) = m^{\Omega(1)}$, then for any constant t , $G^{(t)}$ is a cryptographic pseudorandom generator. Note where your proof fails for fully explicit pseudorandom generators against time m^c for a fixed constant c .

Problem 7.12 (Local List Decoding the Hadamard Code). For a function $f : \mathbb{Z}_2^m \rightarrow \mathbb{Z}_2$, a parameterized subspace $x + V$ of \mathbb{Z}_2^m of dimension d is given by a linear map $V : \mathbb{Z}_2^d \rightarrow \mathbb{Z}_2^m$ and a shift $x \in \mathbb{Z}_2^m$. (We do not require that the map V be full rank.) We write V for $0 + V$. For a function $f : \mathbb{Z}_2^m \rightarrow \mathbb{Z}_2$, we define $f|_{x+V} : \mathbb{Z}_2^d \rightarrow \mathbb{Z}_2^m$ by $f|_{x+V}(y) = f(x + V(y))$.

- (1) Let $c : \mathbb{Z}_2^m \rightarrow \mathbb{Z}_2$ be a codeword in the Hadamard code (i.e. a linear function), $r : \mathbb{Z}_2^m \rightarrow \mathbb{Z}_2$ a received word, V a parameterized subspace of \mathbb{Z}_2^m of dimension d , and $x \in \mathbb{Z}_2^m$. Show that if $d_H(r|_{x+V}, c|_{x+V}) < 1/2$, then $c(x)$ can be computed from x , V , $c|_V$, and oracle access to r in time $\text{poly}(m, 2^d)$ with $2^d - 1$ queries to r .
 - (2) Show that for every $m \in \mathbb{N}$ and $\varepsilon > 0$, the Hadamard code of dimension m has a $(1/2 - \varepsilon)$ local list-decoding algorithm $(\text{Dec}_1, \text{Dec}_2)$ in which both Dec_1 and Dec_2 run in time $\text{poly}(m, 1/\varepsilon)$, and the list output by Dec_1 has size $O(1/\varepsilon^2)$. (Hint: consider a random parameterized subspace V of dimension $2 \log(1/\varepsilon) + O(1)$, and how many choices there are for $c|_V$.)
 - (3) Show that Dec_2 can be made to be deterministic and run in time $O(m)$.
-

Problem 7.13 (Hardcore Predicates). A *hardcore predicate* for a one-way function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ is a poly(ℓ)-time computable function $b : \{0, 1\}^\ell \rightarrow \{0, 1\}$ such that for every constant c , every nonuniform algorithm A running in time ℓ^c , we have:

$$\Pr[A(f(U_\ell)) = b(U_\ell)] \leq \frac{1}{2} + \frac{1}{\ell^c},$$

for all sufficiently large ℓ . Thus, while the one-wayness of f only guarantees that it is hard to compute *all* the bits of f 's input from its output, b specifies a particular bit of information about the input that is very hard to compute (one can't do noticeably better than random guessing).

- (1) Let $\text{Enc} : \{0, 1\}^\ell \rightarrow \{0, 1\}^{\hat{L}}$ be a code such that given $x \in \{0, 1\}^\ell$ and $y \in [\hat{L}]$, $\text{Enc}(x)_y$ can be computed in time $\text{poly}(\ell)$. Suppose that for every constant c and all sufficiently large ℓ , Enc has a $(1/2 - 1/\ell^c)$ local list-decoding algorithm $(\text{Dec}_1, \text{Dec}_2)$ in which both Dec_1 and Dec_2 run in time $\text{poly}(\ell)$. Prove that if $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ is a one-way function, then $b(x, y) = \text{Enc}(x)_y$ is a hardcore predicate for the one-way function $f'(x, y) = (f(x), y)$.
- (2) Show that if $b : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is a hardcore predicate for a one-way permutation $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$, then for every $m = \text{poly}(\ell)$, the following function $G : \{0, 1\}^\ell \rightarrow \{0, 1\}^m$ is a cryptographic pseudorandom generator:

$$G(x) = (b(x), b(f(x)), b(f(f(x))), \dots, b(f^{(m-1)}(x))).$$

(Hint: show that G is “previous-bit unpredictable”.)

- (3) Using Problem 7.12, deduce that if $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ is a one-way permutation, then for every $m = \text{poly}(\ell)$, the following is a cryptographic pseudorandom generator:

$$G_m(x, r) = (\langle x, r \rangle, \langle f(x), r \rangle, \langle f(f(x)), r \rangle, \dots, \langle f^{(m-1)}(x), r \rangle).$$

Problem 7.14 (PRGs from 1–1 One-Way Functions). A random variable X has (t, ε) pseudoentropy at least k if it is (t, ε) indistinguishable from some random variable of min-entropy at least k .

- (1) Suppose that X has (t, ε) pseudoentropy at least k and that $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε') -extractor computable in time t' . Show that $\text{Ext}(X, U_d)$ is an $(t-t', \varepsilon+\varepsilon')$ indistinguishable from U_m .
- (2) Let $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^{\ell'}$ be a one-to-one one-way function (not necessarily length-preserving) and $b : \{0, 1\}^{\ell'} \rightarrow \{0, 1\}$ a hardcore predicate for f (see Problem 7.13). Show that for every constant c and all sufficiently large ℓ , the random variable $f(U_\ell)b(U_\ell)$ has $(\ell^c, 1/\ell^c)$ pseudoentropy at least $\ell + 1$.
- (3) (*) Show how to construct a cryptographic pseudorandom generator from any one-to-one one-way function. (Any seed length $\ell(m) < m$ is fine.)

7.9 Chapter Notes and References

Other surveys on pseudorandom generators and derandomization include [Mil2, Imp2, Kab, Gol8].

Descriptions of classical constructions of pseudorandom generators (e.g. linear congruential generators) and the batteries of statistical tests that are used to evaluate them can be found in [Knu, RSN⁺]. Linear congruential generators and variants were shown to be cryptographically insecure (e.g. not satisfy Definition 7.9) in [Boy, FHK⁺, Ste2, Kra, BGM]. Current standards for pseudorandom generation in practice can be found in [BK].

The modern approach to pseudorandomness described in this chapter grew out of the complexity-theoretic approach to cryptography initiated by Diffie and Hellman [DH] (who introduced the concept of one-way functions, among other things). Shamir [Sha2] constructed a generator achieving a weak form of unpredictability based on the conjectured one-wayness of the RSA function [RSA]. (Shamir's generator outputs a sequence of long strings, such that none of the string can be predicted from the others, except with negligible probability, but individual bits may be easily predictable.) Blum and Micali [BM] proposed the criterion of next-bit unpredictability (Definition 7.15) and constructed a generator satisfying it based on the conjectured hardness of the Discrete Logarithm Problem. Yao [Yao] gave the now-standard definition of pseudorandomness (Definition 7.3) based on the notion of computational indistinguishability introduced in the earlier work of Goldwasser and Micali [GM] (which also introduced hybrid arguments). Yao also proved the equivalence of pseudorandomness and next-bit unpredictability (Proposition 7.16), and showed how to construct a cryptographic pseudorandom generator from any one-way permutation. The construction described in Section 7.2 and Problems 7.12 and 7.13 uses the hardcore predicate from the later work of Goldreich and Levin [GL]. The construction of a pseudorandom generator from an arbitrary one-way function (Theorem 7.11) is due to Håstad, Impagliazzo, Levin, and Luby [HILL]. The most efficient (and simplest) construction of pseudorandom generators from general one-way functions to date is in [HRV, VZ]. Goldreich, Goldwasser, and Micali [GGM] defined and constructed pseudorandom functions, and illustrated their applicability in cryptography. The application of pseudorandom functions to learning theory is from [Val3], and their application to

circuit lower bounds is from [RR]. For more about cryptographic pseudorandom generators, pseudorandom functions, and their applications in cryptography, see the text by Goldreich [Gol3].

Yao [Yao] demonstrated the applicability of pseudorandom generators to derandomization, noting in particular that cryptographic pseudorandom generators imply that $\mathbf{BPP} \subset \mathbf{SUBEXP}$, and that one can obtain even $\mathbf{BPP} \subset \tilde{\mathbf{P}}$ under stronger intractability assumptions. Nisan and Wigderson [NW] observed that derandomization only requires a mildly explicit pseudorandom generator, and showed how to construct such generators based on the average-case hardness of \mathbf{E} (Theorem 7.24). A variant of Open Problem 7.25 was posed in [HVV], who showed that it also would imply stronger results on hardness amplification; some partial negative results can be found in [IW1, RRV].

The instantiation of the Nisan–Wigderson pseudorandom generator that uses the parity function to fool constant-depth circuits (Theorem 7.29) is from the earlier work of Nisan [Nis1]. (The average-case hardness of parity against constant-depth circuits stated in Theorem 7.27 is due Boppana and Håstad [Hås].) The first unconditional pseudorandom generator against constant-depth circuits was due to Ajtai and Wigderson [AW] and had seed length $\ell(m) = m^\epsilon$ (compared to $\text{polylog}(m)$ in Nisan’s generator). Recently, Braverman [Bra] proved that any $\text{polylog}(m)$ -wise independent distribution fools \mathbf{AC}^0 , providing a different way to obtain polylogarithmic seed length and resolving a conjecture of Linial and Nisan [LN]. The notion of strong pseudorandom generators (a.k.a. seed-extending pseudorandom generators) and the average-case derandomization of \mathbf{AC}^0 (Problem 7.5) are from [Sha4, KvMS]. Superpolynomial circuit lower bounds for $\mathbf{AC}^0[2]$ were given by [Raz, Smo]. Viola [Vio2] constructed pseudorandom generators with superpolynomial stretch for $\mathbf{AC}^0[2]$ circuits that are restricted to have a logarithmic number of parity gates.

Detailed surveys on locally decodable codes and their applications in theoretical computer science are given by Trevisan [Tre3] and Yekhanin [Yek2]. The notion grew out of several different lines of work, and it took a couple of years before a precise definition of locally decodable codes was formulated. The work of Goldreich and Levin [GL] on hardcore predicates of one-way permutations implicitly provided a local list-decoding algorithm for the Hadamard code. (See Problems 7.12 and 7.13.) Working on the problem of “instance hiding” introduced in [AFK], Beaver and Feigenbaum [BF] constructed a protocol based on Shamir’s “secret sharing” [Sha1] that effectively amounts to using the local decoding algorithm for the Reed–Muller code (Algorithm 7.43) with the multilinear extension (Lemma 7.47). Blum, Luby, and Rubinfeld [BLR] and Lipton [Lip] introduced the concept of self-correctors for functions, which allow a one to convert a program that correctly computes a function on most inputs to one that correctly computes the function on all inputs.⁹ Both papers gave self-correctors for group homomorphisms, which, when applied to homomorphisms from \mathbb{Z}_2^n to \mathbb{Z}_2 , can be interpreted as a local corrector for the Hadamard code (Proposition 7.40). Lipton [Lip] observed that the techniques of Beaver and Feigenbaum [BF] yield a self-corrector for multivariate polynomials, which, as mentioned above, can be interpreted as a local corrector for the Reed–Muller code. Lipton pointed out that it is interesting to apply these self-correctors to presumably intractable functions, such as the PERMANENT (known to be $\#\mathbf{P}$ -complete [Val2]),

⁹Blum, Luby, and Rubinfeld [BLR] also defined and constructed self-testers for functions, which allow one to efficiently determine whether a program does indeed compute a function correctly on most inputs before attempting to use self-correction. Together a self-tester and self-corrector yield a “program checker” in the sense of [BK]. The study of self-testers gave rise to the notion of *locally testable codes*, which are intimately related to probabilistically checkable proofs [BFL, BFLS], and to the notion of *property testing* [RS, GGR, Ron], which is an area within sublinear-time algorithms.)

and soon it was realized that they could also be applied to complete problems for other classes by taking the multilinear extension [BFL]. Babai, Fortnow, Nisan, and Wigderson [BFNW] used these results to construct pseudorandom generators from the worst-case hardness of **EXP** (or **E**, due to Problem 7.2), and thereby obtain subexponential-time or quasipolynomial-time simulations of **BPP** under appropriate worst-case assumptions (Corollary 7.64, Parts 2 and 3). All of these works also used the terminology of random self-reducibility, which had been present in the cryptography literature for a while [AL], and was known to imply worst-case/average-case connections. Understanding the relationship between the worst-case and average-case complexity of **NP** (rather than “high” classes like **EXP**) is an important area of research; see the survey [BT].

Self-correctors for multivariate polynomials that can handle a constant fraction of errors (as in Theorem 7.42) and fraction of errors approaching $1/2$ (as in Problem 7.7) were given by Gemmell et al. [GLR⁺] and Gemmell and Sudan [GS], respectively. Babai, Fortnow, Levin, and Szegedy [BFLS] reformulated these results as providing error-correcting codes with efficient local decoding (and “local testing”) algorithms. Katz and Trevisan [KT] focused attention on the exact query complexity of locally decodable codes (separately from computation time), and proved that locally decodable codes cannot simultaneously have the rate, distance, and query complexity all be constants independent of the message length. Constructions of 3-query locally decodable codes with subexponential blocklength were recently given by Yekhanin [Yek1] and Efremenko [Efr]. Private Information Retrieval (Problem 7.6) was introduced by Chor, Goldreich, Kushilevitz, and Sudan [CGKS]. Katz and Trevisan [KT] introduced the notion of smooth codes and showed their close relation to both private information retrieval and locally decodable codes (Problem 7.6). Recently, Saraf, Kopparty, and Yekhanin [KSY] constructed the first locally decodable codes with sublinear-time decoding and rate larger $1/2$.

Techniques for Hardness Amplification (namely, the Direct Product Theorem and XOR Lemma) were first described in oral presentations of Yao’s paper [Yao]. Since then, these results have been strengthened and generalized in a number of ways. See the survey [GNW] and Section 8.2.3. The first local list-decoder for Reed–Muller codes was given by Arora and Sudan [AS2] (stated in the language of program self-correctors). The one in Theorem 7.56 is due to Sudan, Trevisan, and Vadhan [STV], who also gave a general definition of locally list-decodable codes (inspired by a list-decoding analogue of program self-correctors defined by Ar et al. [ALRS]) and explicitly proved Theorems 7.60, 7.61, and 7.62.

The result that **BPP** = **P** if **E** has a function of nonuniform worst-case hardness $s(\ell) = 2^{\Omega(\ell)}$ (Corollary 7.64, Part 1) is from the earlier work of Impagliazzo and Wigderson [IW2], who used derandomized versions of the XOR Lemma to obtain sufficient average-case hardness for use in the Nisan–Wigderson pseudorandom generator. An optimal construction of pseudorandom generators from worst-case hard functions, with seed length $d(m) = O(s^{-1}(\text{poly}(m)))$ (cf., Theorem 7.63), was given by Shaltiel and Umans [SU1, Uma].

For more background on **AM**, see the Notes and References of Chapter 2. The first evidence that **AM** = **NP** was given by Arvind and Köbler [AK], who showed that one can use the Nisan–Wigderson generator with a function that is $(2^{\Omega(\ell)}, 1/2 - 1/2^{\Omega(\ell)})$ -hard for nondeterministic circuits. Klivans and van Melkebeek [KvM] observed that the Impagliazzo–Wigderson pseudorandom generator construction is “black box” and used this to show that **AM** can be derandomized using functions that are worst-case hard for circuits with an **NP** oracle (Theorem 7.68). Subsequent work showed that one only needs worst-case hardness against a nonuniform analogue of

$\mathbf{NP} \cap \mathbf{co-NP}$ [MV, SU1, SU2].

Trevisan [Tre1] showed that black-box pseudorandom generator constructions yield randomness extractors, and thereby obtained the extractor construction of Theorem 7.73. This surprising connection between complexity-theoretic pseudorandomness and information-theoretic pseudorandomness sparked much subsequent work, from which the unified theory presented in this survey emerged. The fact that black-box hardness amplifiers are a form of locally list-decodable codes was explicitly stated (and used to deduce lower bounds on advice length) in [TV2]. The use of black-box constructions to classify and separate cryptographic primitives was pioneered by Impagliazzo and Rudich [IR]; see also [RTV1, Reil].

Problem 7.1 (PRGs imply hard functions) is from [NW]. Problem 7.2 is a special case of the technique called “translation” or “padding” in complexity theory. Problem 7.4 (Deterministic Approximate Counting) is from [NW]. The fastest known deterministic algorithms for approximately counting the number of satisfying assignments to a DNF formula are from [LVW] and [GMR2] (depending on whether the approximation is relative or additive, and the magnitude of the error). The fact that hitting set generators imply $\mathbf{BPP} = \mathbf{P}$ (Problem 7.8) was first proven by Andreev, Clementi, and Rolim [ACR]; for a more direct proof, see [GVW]. Problem 7.9 (that PRGs vs. uniform algorithms imply average-case derandomization) is from [IW3]. Goldreich [Gol9] showed that PRGs are necessary for derandomization (Problem 7.10). The result that one-to-one one-way functions imply pseudorandom generators is due to Goldreich, Krawczyk, and Luby [GKL]; the proof in Problem 7.14 is from [HILL].

For more on Kolmogorov complexity, see [LV2]. In recent years, connections have been found between Kolmogorov complexity and derandomization; see [ABK⁺]. The tighter equivalence between circuit size and nonuniform computation time mentioned after Definition 7.1 is due to Pippenger and Fischer [PF]. The $5n - o(n)$ lower bound on circuit size is due to Iwama, Lachish, Morizumi, and Raz [LR, IM]. The fact that single-sample indistinguishability against uniform algorithms does not imply multiple-sample indistinguishability unless we make additional assumptions such as efficient samplability (in contrast to Proposition 7.14), is due to Goldreich and Meyer [GM]. (See also [GS].)

8

Conclusions

8.1 A Unified Theory of Pseudorandomness

In the course of this survey, we have seen that a wide variety of “pseudorandom objects” are very closely related, and indeed almost equivalent when viewed appropriately: list-decodable codes, averaging samplers, expander graphs, randomness extractors, (black-box) pseudorandom generator constructions, and hardness amplifiers. The power of these connections comes from the fact they allow us to translate intuitions and techniques that are very natural for one of these objects to others where they have been might be difficult to discover. Indeed, we have seen several examples of this, such the use of randomness extractors to construct near-optimal samplers (Problem 6.5), the use of Parvaresh–Vardy codes to construct bipartite expanders (Thm. 5.35) and lossless condensers (Thm. 6.34), and the use of the Nisan–Wigderson pseudorandom generator construction to construct extractors (Thm. 7.73).

Table 8.1 reviews the connections between these objects, by putting them in our list-decoding framework. Recall that we can view each of the objects as a function $\Gamma : [N] \times [D] \rightarrow [M]$ and can characterize their properties by bounding the sizes of sets of the form $\text{LIST}_\Gamma(T, \alpha) = \{x : \Pr_y[\Gamma(x, y) \in T] > \alpha\}$, where $T \subset [M]$. (When $\alpha = 1$, we change the inequality to an equality.) Other objects we have encountered, such as dispersers (Def. 6.19), black-box hitting-set generator constructions (Problem 7.8), and lossy condensers (Def. 6.32) can also be cast in the framework, but we leave this as an exercise.

In addition to illustrating the connections between the objects, Table 8.1 also brings out their differences. In some cases $\alpha = \mu(T) + \varepsilon$, and in other cases $\alpha = 1$. In some cases (samplers, extractors, PRG constructions), we consider all subsets $T \subset [M]$, and in other cases (codes, hardness amplifiers, expanders, condensers), we only consider T that are small and possibly have some additional structure, like being the graph of a received word. (Typically, the only way the graph structure is used is to determine that the size of T is D .) In some cases, we want efficient algorithms to construct $\text{LIST}(T, \alpha)$ (even polylogarithmic-time local decoding algorithms), and in others, all we care about is its size.

Object	$\Gamma(x, y)$	List-Decoding Problem	Decoding Time	Standard Parameters	Prop.
$(1 - 1/q - \epsilon, K)$ list-decodable codes	$(y, \text{Enc}(x)_y)$	$T = \{(y, r_y)\} \Rightarrow$ $ \text{LIST}(T, \mu(T) + \epsilon) \leq K$	$\text{poly}(n, 1/\epsilon)$	$q = O(1), \epsilon = \Omega(1),$ $M = Dq = O(n), K = \text{poly}(n)$	5.29
(t, k, ϵ) black-box hardness amplifiers	$(y, \text{Amp}^f(y))$	$T = \{(y, r_y)\} \Rightarrow$ $ \text{LIST}(T, \mu(T) + \epsilon) \leq K$ $ \text{LIST}(T, \mu(T) + \epsilon) \leq K$	$t = \text{poly}(\log n, 1/\epsilon)$ (local w/ advice)	$d = O(\log n), q = 2, M = Dq = \text{poly}(n),$ $k = t = \text{poly}(\log n, 1/\epsilon) \leq n^\gamma$	7.74
(δ, ϵ) averaging samplers	$\text{Samp}(x)_y$	$ \text{LIST}(T, \mu(T) + \epsilon) \leq K$	don't care	$K = \delta N, D = \text{poly}(1/\epsilon, \log(1/\delta)),$ $n = O(m + \log(1/\delta))$	5.30
(k, ϵ) extractors	$\text{Ext}(x, y)$	$ \text{LIST}(T, \mu(T) + \epsilon) \lesssim K$	don't care	$d = O(\log(n/\epsilon)),$ $k = O(m) \in [\text{poly}\log(n), \gamma m]$	6.23
(t, k, ϵ) black-box PRGs	$G^x(y)$	$ \text{LIST}(T, \mu(T) + \epsilon) \leq K$	$t = \text{poly}(m)$ (local w/ advice)	$\epsilon = 1/m,$ $k = t = \text{poly}(m) \in [\text{poly}\log(n), n^\gamma]$	7.74
$(= K, A)$ expanders	y^{th} nbr of x	$ T < AK \Rightarrow$ $ \text{LIST}(T, 1) < K$	don't care	$D = O(1), A = 1 + \Omega(1),$ $K = N/2, M = N$	5.33
(δ, ϵ) hitting samplers	$\text{Samp}(x)_y$	$ T < (1 - \epsilon)M \Rightarrow$ $ \text{LIST}(T, \mu(T) + \epsilon) \leq K$	don't care	$K = \delta N, D = \text{poly}(1/\epsilon, \log(1/\delta)),$ $n = O(m + \log(1/\delta))$	Prob. 4.7
$k \rightarrow_\epsilon k + d$ (lossless) condensers	$\text{Con}(x, y)$	$ T < (1 - \epsilon)DK \Rightarrow$ $ \text{LIST}(T, 1) < K$	don't care	$d = O(\log(m/\epsilon)),$ $k = O(m) \in [\text{poly}\log(n), \gamma m]$	6.33

Table 8.1 Capturing Pseudorandom Objects $\Gamma : [N] \times [D] \rightarrow [M]$ by bounding sizes of $\text{LIST}_\Gamma(T, \alpha) = \{x : \text{Pr}_y[\Gamma(x, y) \in T] > \alpha\}$ for $T \subset [M]$. The parameter γ denotes a arbitrarily small positive constant. As usual, $N = 2^n, M = 2^m, D = 2^d$, and $K = 2^k$.

The most common parameter ranges also vary quite widely among the objects. The typical relation between M and N ranges from logarithmic (codes) to polylogarithmic (hardness amplifiers) to a constant power (samplers) to a constant multiplicative factor (expanders). (Extractors, PRG constructions, and condensers all consider the full range between polylogarithmic and a constant power.) The typical value of D ranges from being a constant independent of N (expanders), to being logarithmic in N (codes), to being polylogarithmic in N (extractors, condensers, PRG constructions, hardness amplifiers). Despite these differences, we have seen that the connections are nevertheless quite powerful, and ideas or techniques developed for one of these objects are often useful for the others.

The unified framework of Table 8.1 opens up a tantalizing possibility that we can also have a unified *construction* of all of these objects. In what follows, we will argue that this is possible ignoring explicitness/efficiency considerations (including local list-decoding). The bounds will be stated in terms of the tails of the Binomial distribution:

Definition 8.1 (tails of Binomial distributions). For $t \in \mathbb{N}$, $\mu, \alpha \in (0, 1)$, define

$$\text{Bin}(t, \mu, \alpha) = \Pr \left[\frac{1}{t} \sum_{i=1}^t X_i > \alpha \right],$$

where the X_i 's are iid $\{0, 1\}$ -valued random variables such that $\Pr[X_i = 1] = \mu$, and define

$$\text{Bin}(t, \mu, 1) = \Pr \left[\frac{1}{t} \sum_{i=1}^t X_i = 1 \right] = \mu^t.$$

With this definition, we can give the following unified construction:

Theorem 8.2 (nonconstructive unified pseudorandom object). For every $N, M, D \in \mathbb{N}$, there exists a function $\Gamma : [N] \times [D] \rightarrow [M]$ such that for every $T \subset [M]$ and every $\alpha \in (0, 1]$, we have

$$|\text{LIST}_\Gamma(T, \alpha)| \leq \min \left\{ K \in \mathbb{N} : \binom{N}{K} \cdot \binom{M}{|T|} \cdot \text{Bin}(D, \mu(T), \alpha)^K \leq \frac{1}{4K^2|T|^2} \right\}.$$

We leave the proof of Theorem 8.2 as an exercise and instead show how many of the nonconstructive bounds we've seen for various pseudorandom objects would simultaneously hold for the Γ of Theorem 8.2 by setting parameters appropriately:

- The nonconstructive bounds for expanders (Theorem 4.4) can be obtained by setting $M = N$, $\alpha = 1$, $\mu = \mu(T)$, $K = |T|/(D - 2)$ (ignoring round-off errors), and noting that

$$\begin{aligned} \binom{N}{K} \cdot \binom{N}{|T|} \cdot \text{Bin}(D, \mu, 1)^K &\leq \left(\frac{Ne}{K} \right)^K \cdot \left(\frac{Ne}{|T|} \right)^{|T|} \cdot \mu^{DK} \\ &= \left(\frac{(D-2) \cdot e}{\mu} \right)^K \cdot \left(\frac{e}{\mu} \right)^{(D-2) \cdot K} \cdot \mu^{D \cdot K} \\ &= ((D-2)e^{D-1}\mu)^K \\ &< 1/(4K^2|T|^2), \end{aligned}$$

provided μ is sufficiently small. Thus, by Proposition 5.33, Γ defines a $(\gamma N, (D - 2))$ vertex expander for a sufficiently small γ .

More generally, we see that for any constants $\gamma, A, D > 0$ and sufficiently large $N = M$, Γ will be a $(\gamma N, A)$ expander provided that $D > (H(\gamma) + H(\gamma A))/\gamma \log(1/\gamma A)$, where $H(x) = x \log(1/x) + (1 - x) \log(1/(1 - x))$ is the binary entropy function.¹ (To see this, use the fact that $\binom{N}{\beta N} = 2^{(H(\beta) + o(1)) \cdot N}$ as $N \rightarrow \infty$.)

- For the nonconstructive bound on extractors (Theorem 6.14) and averaging samplers, we set $\alpha = \mu(T) + \varepsilon$ and use the bounds $\binom{N}{K} \leq (Ne/K)^K$, $\binom{M}{T} \leq 2^M$, and $\text{Bin}(D, \mu(T), \mu(T) + \varepsilon) \leq \exp(-\Omega(D\varepsilon^2))$ (by the Chernoff Bound of Theorem 2.21). Then the condition of Theorem 8.2 is satisfied provided that $D \geq (c/\varepsilon^2) \cdot \log(N/K)$ and $M \leq \varepsilon^2 KD/c$ for a sufficiently large constant c .

To get the “strong” forms of pseudorandom objects, where $\Gamma(x, y) = (y, \Gamma'(x, y))$ for some function Γ' , we need to consider the tails of sums of independent, but not necessarily identically distributed, Binomial random variables:

Definition 8.3 (tails of Poisson Binomial distributions). For $t, \mathbb{N}, \mu, \alpha \in (0, 1)$, define

$$\text{PBin}(t, \mu, \alpha, \eta) = \sup_{(Y_1, \dots, Y_t)} \Pr \left[\frac{1}{t} \sum_{i=1}^t Y_i > \alpha \right],$$

and

$$\text{PBin}(t, \mu, 1, \eta) = \sup_{(Y_1, \dots, Y_t)} \Pr \left[\frac{1}{t} \sum_{i=1}^t Y_i = 1 \right] = \text{Bin}(t, \mu, 1),$$

where the suprema are taken over all sequences of independent Bernoulli random variables Y_1, \dots, Y_t such that $(1/t) \cdot \sum_{i=1}^t \mathbb{E}[Y_i] \leq \mu$, and additionally $\mathbb{E}[Y_i]$ is an integer multiple of η for every i .

Theorem 8.4 (nonconstructive unified pseudorandom object). For every $N, q, D \in \mathbb{N}$, there exists a function $\Gamma' : [N] \times [D] \rightarrow [q]$ such that if we define $\Gamma(x, y) = (y, \Gamma'(x, y))$, then for every $T \subset [D] \times [q]$ and every $\alpha \in (0, 1]$, we have

$$|\text{LIST}_{\Gamma}(T, \alpha)| \leq \min \left\{ K \in \mathbb{N} : \binom{N}{K} \cdot \binom{Dq}{|T|} \cdot \text{PBin}(D, \mu(T), \alpha, 1/q)^K \leq \frac{1}{4K^2|T|^2} \right\}.$$

- Nonconstructive bounds for strong vertex expanders, hitting samplers, and lossless condensers (matching the bounds for the non-strong ones) follow by noting that $\text{PBin}(n, \mu, 1, \eta) \leq \text{Bin}(n, \mu, 1)$. Indeed, writing μ_i for the expectation of Bernoulli random variable Y_i , then

$$\Pr \left[\frac{1}{t} \sum_{i=1}^t Y_i = 1 \right] = \prod_{i=1}^t \mu_i \leq \left(\frac{1}{t} \sum_{i=1}^t \mu_i \right)^t \leq \mu^t,$$

¹The discussion after Theorem 4.4 states a slightly stronger result, only requiring $D > (H(\gamma) + H(\gamma A))/(H(\gamma) - \gamma A H(1/\gamma))$, but this is based on choosing a graph that is the union of D random perfect matchings, which does indeed have slightly better expansion properties than the model we are using now, of choosing D random neighbors for each left vertex.

where the second-to-last inequality follows from the Arithmetic Mean–Geometric Mean Inequality.

- Nonconstructive bounds for strong extractors and averaging samplers (matching the bounds for the non-strong ones) follow because the Chernoff Bound (Theorem 2.21) also applies to non-identical random variables: $\text{PBin}(D, \mu(T), \mu(T) + \varepsilon, \eta) \leq \exp(-\Omega(D\varepsilon^2))$ for every $\eta > 0$.
- Nonconstructive bounds for list-decodable codes (Thm. 5.8) follow because the list-decodability of codes involves taking T to be the graph of a received word (cf. Proposition 5.29) and hence $\mu(T) = 1/q$, and $\text{PBin}(D, 1/q, 1 - \delta, 1/q) = q^{H_q(\delta, D) \cdot D} / q^D$. To see the latter, let $\mathbb{E}[Y_i] = m_i/q$ for $m_i \in \mathbb{N}$ such that $\sum_i m_i \leq D$. so Y_i can be viewed as the indicator for the event the $R(i) \in \{1, \dots, m_i\}$ for a uniformly random received word $R : [D] \rightarrow [q]$, and $\Pr[\sum_i Y_i > 1 - \delta]$ is the probability that the graph $G(R) = \{(i, r(i)) : i \in [D]\}$ of the received word R intersects the set $T' = \{(i, j) : i \in [D], j \leq m_i\}$ in more than $(1 - \delta)D$ positions. For a set $S \subset T'$, the probability that $G(R) \cap T' = S$ is at most $(1/q)^{|S|} \cdot (1 - 1/q)^{D - |S|}$. (This is exactly the probability in case S contains only pairs (i, j) with D distinct values of i , otherwise the probability is zero.) Thus,

$$\begin{aligned} \Pr \left[\sum_i Y_i > 1 - \delta \right] &\leq \sum_{s > (1 - \delta)D} \binom{\sum_i m_i}{s} (1/q)^s \cdot (1 - 1/q)^{D - s} \\ &\leq \frac{1}{q^D} \cdot \sum_{s > (1 - \delta)D} \binom{D}{s} (q - 1)^s \\ &= \frac{q^{H_q(\delta, D) \cdot D}}{q^D}. \end{aligned}$$

Now, given this bound on $\text{PBin}(\cdot)$ and Proposition 5.29, Theorem 8.4 gives us a (δ, K) list-decodable code $\text{Enc} : [N] \rightarrow [q]^D$ provided that $\binom{N}{K} \cdot \binom{Dq}{D} \cdot \left(\frac{q^{H_q(\delta, D) \cdot D}}{q^D} \right)^K \leq \frac{1}{4K^2 D^2}$. Using the bounds $\binom{N}{K} \leq N^K$ and $\binom{Dq}{D} \leq (eq)^D$, we see that we get a list-decodable code provided that the rate $\rho = \log N / D \log q$ is smaller than $H_q(\delta, D) - O(1/K) - O((\log KD) / (KD \log q))$. This matches the nonconstructive bound for list-decodable codes (Thm. 5.8) up to the dependence of the error term on the list size K .

Given the above, a natural long-term goal for the theory of pseudorandomness is the following:

Open Problem 8.5. Can we give explicit constructions of functions Γ that nearly match the bounds of Theorems 8.2 and 8.4? Can this be done with efficient (local) list-decoding algorithms?

8.2 Other Topics in Pseudorandomness

In this section, we survey some additional topics in the theory of pseudorandomness that we did not cover in depth, each of which could merit at least an entire chapter on its own. Sections 7.2 and 7.9 contained a detailed survey of cryptographic pseudorandomness, so we do not discuss it again here.

8.2.1 Pseudorandomness for Space-Bounded Computation

As discussed in Section 4.4, unlike for **BPP** and **RP**, we do know very nontrivial derandomizations **RL** (and **BPL**). These derandomizations are obtained constructions of pseudorandom generators $G : \{0, 1\}^d \rightarrow \{0, 1\}^m$ such that no randomized $(\log m)$ -space algorithm can distinguish $G(U_d)$ from U_m . In order to get derandomizations that are correct on every input x , we require pseudorandom generators that fool *nonuniform* space-bounded algorithms. On the other hand, since randomized space-bounded algorithms get their random bits as a stream of coin tosses, we only need to fool space-bounded distinguishers that read each of their input bits once, in order. Thus, we consider distinguishers that are (*oblivious, read-once*) *branching programs*, which maintain a state $s_i \in [w]$ after reading i input bits, and determine the next state $s_{i+1} \in [w]$ as a function of s_i and the $(i + 1)$ 'st input bit. The number w of available states at each time step is called the *width* of the branching program, and corresponds to a space bound of $\log w$ bits. A generator $G : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$ that is computable in space $O(d(m))$ and such that $G(U_{d(m)})$ cannot be distinguished from U_m by oblivious, read-once branching programs of width m implies that $\mathbf{BPL} \subset \mathbf{DSPACE}(O(d(\text{poly}(m))))$.

The fact that pseudorandom generators imply lower bounds (Problem 7.1) applies to this context too, but fortunately we do know exponential width lower bounds for *oblivious, read-once* branching programs (e.g. computing the inner-product modulo 2 of two ℓ -bit strings requires width $2^{\Omega(\ell)}$). On the other hand, we cannot simply plug such lower bounds into the Nisan–Wigderson generator construction (Theorem 7.24), because the reductions used do not preserve the read-once property (and we do not know superpolynomial lower-bounds for branching programs that can read each input bit many times).

Nevertheless, a series of papers starting with Ajtai, Komlós, and Szemerédi [AKS2] have given *unconditional* pseudorandom generators for space-bounded computation. The pseudorandom generator of Nisan [Nis2] uses a seed of length $O(\log^2 m)$ to produce m bits that are pseudorandom to oblivious, read-once branching programs of width m . At first, this only seems to imply that $\mathbf{RL} \subset \mathbf{L}^2$, which already follows from Savitch's Theorem [Sav] that $\mathbf{NL} \subset \mathbf{L}^2$. Nevertheless, Nisan's generator and its additional properties has been used in more sophisticated ways to obtain highly nontrivial derandomizations of **RL**. Specifically, Nisan [Nis3] used it to show that every problem in **RL** can be solved simultaneously in polynomial time and $O(\log^2 n)$ space, and Saks and Zhou [SZ1] used it to prove that $\mathbf{RL} \subset \mathbf{L}^{3/2}$. Another important pseudorandom generator for space-bounded computation is that of Nisan and Zuckerman [NZ], which uses a seed of length $O(\log m)$ to produce $\log^k m$ bits that are pseudorandom to oblivious, read-once branching programs of width m . None of these results have been improved in nearly two decades.

However, substantially better generators have been constructed for restricted classes of oblivious read-once branching programs. Specifically, there are pseudorandom generators or hitting-set generators (see Problem 7.8) stretching a seed of length $\tilde{O}(\log m)$ to m bits that fool combinatorial rectangles (which check membership in a “rectangle” $S_1 \times S_2 \times \cdots \times S_{m/\log m}$, where each $S_i \subset \{0, 1\}^{\log m}$) [EGL⁺, LLSZ, ASWZ, Lu, GMR⁺], branching programs of width 2 and 3 [SZ2, BDVY, SZ, GMR⁺], constant-width *regular* branching programs (where the transition function at each layer is regular) [BRRY, BV], and constant-width *permutation* branching programs (where each input string induces a permutation of the states at each layer) [RV, KNP, De, Ste1]. However, the following remains open:

Open Problem 8.6. Is there an explicit pseudorandom generator $G : \{0, 1\}^{o(\log^2 m)} \rightarrow \{0, 1\}^m$ whose output distribution is pseudorandom to oblivious, read-once branching programs of width 4?

8.2.2 Derandomization vs. Lower Bounds

Derandomization from Uniform Assumptions. The construction of pseudorandom generators we have seen (Theorem 7.63) requires nonuniform circuit lower bounds for functions in \mathbf{E} , and it is of interest to find constructions that only require uniform hardness assumptions. One way to obtain such results is to show that uniform hardness assumptions imply nonuniform assumptions. Indeed, the Karp–Lipton Theorems in complexity theory [KL] show that certain strong uniform lower bounds imply nonuniform lower bounds. For example, if an \mathbf{EXP} -complete problem cannot be solved by a uniform algorithm in the second level of the polynomial-time hierarchy (like \mathbf{NP} but with two nondeterministic quantifiers, cf. Section 3.3), then it also cannot be solved by polynomial-sized circuits. Subsequent strengthenings of the Karp–Lipton Theorem (based on the theory of interactive proofs) show that if $\mathbf{EXP} \neq \mathbf{MA}$ (where \mathbf{MA} is like \mathbf{NP} but allows probabilistic verification of witnesses), then $\mathbf{EXP} \not\subseteq \mathbf{P/poly}$ [BFL]; consequently one gets pseudorandom generators with arbitrary polynomial stretch (secure for infinitely many input lengths) under the assumption $\mathbf{EXP} \neq \mathbf{MA}$ [BFNW].

Assuming the Karp–Lipton Theorem cannot be further improved (e.g. to show $\mathbf{EXP} \not\subseteq \mathbf{P/poly} \Leftrightarrow \mathbf{EXP} \neq \mathbf{BPP}$), from a uniform assumption such as $\mathbf{EXP} \neq \mathbf{BPP}$, we can only hope to construct pseudorandom generators that are secure against *uniform distinguishers* (because pseudorandom generators secure against nonuniform distinguishers imply nonuniform lower bounds, by Problem 7.1). In the context of cryptographic pseudorandom generators, uniform results were typically developed together with or soon after the corresponding nonuniform results. Indeed, analogously to Theorem 7.11, it is known that cryptographic pseudorandom generators that are secure against uniform distinguishers exist if and only if there exist one-way functions that are hard to invert by uniform algorithms [HILL]. For noncryptographic, mildly explicit pseudorandom generators as in Theorem 7.63 and Corollary 7.64, an obstacle is that black-box constructions (Definition 7.65) require nonuniform advice in the reduction. (See discussion at the end of Section 7.7.2. This obstacle is avoided in the case of cryptographic pseudorandom generators, because the appropriate definition of black-box construction from one-way functions gives the reduction oracle access to the one-way function in addition to the distinguisher, since one-way functions are supposed to be efficiently computable by definition.)

Thus, we must turn to *non-black-box* constructions, in which we make more use of the fact that the hard function f is computable in \mathbf{E} and/or the fact that the distinguisher T is computable by an efficient probabilistic algorithm, not just to deduce that G^f is mildly explicit and Red^T is efficient. In fact, f and T need not even be used as oracles; we can make use of the code of the programs computing these functions (e.g. to reduce f to an \mathbf{E} -complete problem). While at first it may seem difficult to take advantage of non-black-box constructions, this was eventually accomplished by Impagliazzo and Wigderson [IW3]. They showed that if $\mathbf{EXP} \neq \mathbf{BPP}$, then there are mildly explicit pseudorandom generators with polynomial stretch that are secure against uniform probabilistic distinguishers (for infinitely many input lengths), and hence \mathbf{BPP} has subexponential-time

average-case derandomizations (by Problem 7.9). (See [TV2] for a precise statement regarding the construction of pseudorandom generators.) This is a uniform analogue of the “low-end” nonuniform result in Corollary 7.64 (Item 3). Analogues for the high-end bounds (Items 1, 2) remain open. For example:

Open Problem 8.7. Does $\mathbf{EXP} \not\subseteq \mathbf{BPSUBEXP}$ imply mildly explicit generators $G : \{0, 1\}^{\text{polylog}(m)} \rightarrow \{0, 1\}^m$ whose input is pseudorandom to every uniform probabilistic algorithm running in time m (for infinitely many m)?

Such a result is known if we replace \mathbf{EXP} by \mathbf{PSPACE} [TV2].

For derandomizing \mathbf{AM} instead of \mathbf{BPP} , both high-end and low-end uniform results have been obtained [GST, SU3]. These results utilize hard functions in \mathbf{E} , unlike the nonuniform results which only require hard functions in $\mathbf{NE} \cap \mathbf{co-NE}$ (cf. Theorem 7.68).

Derandomization Implies Circuit Lower Bounds. Since uniform hardness assumptions and PRGs against uniform distinguishers only seem to imply average-case derandomizations (Problem 7.9), it is tempting to conjecture that *worst-case* derandomizations imply (or are even equivalent to) nonuniform circuit lower bounds. A result of this type was first given implicitly by Buhrman, Fortnow, and Thierauf [BFT] and then explicitly and in stronger form by Impagliazzo, Kabanets, and Wigderson [IKW]. Specifically, these results show that if \mathbf{MA} (like \mathbf{NP} , but with probabilistic verification of witnesses) can be derandomized (e.g. $\mathbf{MA} = \mathbf{NP}$ or even $\mathbf{MA} \subseteq \mathbf{NSUBEXP}$), then $\mathbf{NEXP} \not\subseteq \mathbf{P/poly}$. Derandomization of \mathbf{prBPP} implies derandomization of \mathbf{MA} , so this also implies that if $\mathbf{prBPP} = \mathbf{prP}$ or even $\mathbf{prBPP} \subseteq \mathbf{prSUBEXP}$, then $\mathbf{NEXP} \not\subseteq \mathbf{P/poly}$. This result falls short of giving a converse to Corollary 7.64 (Item 3) because the circuit lower bounds are for \mathbf{NEXP} rather than \mathbf{EXP} . (Corollary 7.64, as well as most of the other derandomization results we’ve seen, apply equally well to \mathbf{prBPP} as to \mathbf{BPP} .) In addition, the result does not give exponential circuit lower bounds even if we assume full derandomization ($\mathbf{prBPP} = \mathbf{prP}$). However, Santhanam [San] shows that $\mathbf{prBPP} = \mathbf{prP}$ implies that for every constant k , there is a language in \mathbf{NP} that does not have circuits of size n^k , which can be viewed as a “scaled down” version of the statement that \mathbf{NE} requires circuits of size $2^{\Omega(n)}$.²

Thus, the following remain open:

Open Problem 8.8. Does $\mathbf{prBPP} = \mathbf{prP}$ imply $\mathbf{E} \not\subseteq \mathbf{P/poly}$ (equivalently $\mathbf{EXP} \not\subseteq \mathbf{P/poly}$, by Problem 7.2)?

Open Problem 8.9. Does $\mathbf{prBPP} = \mathbf{prP}$ imply that \mathbf{NEXP} has a problem requiring nonuniform boolean circuits of size $2^{\ell^{\Omega(1)}}$ on inputs of length ℓ ?

By the result of [IKW] and Corollary 2.31, finding a deterministic polynomial-time algorithm for the \mathbf{prBPP} -complete problem $[\epsilon]$ -APPROX CIRCUIT AVERAGE implies superpolynomial circuit

²Indeed, by a standard “padding” or “translation” argument in complexity theory, if for some constant k , every language in \mathbf{NP} had circuits of size n^k , then every language in \mathbf{NE} would have circuits of size $2^{\Theta(n)}$.

lower bounds for **NEXP**. Unfortunately, we do not know a wide variety of natural problems that are complete for **prBPP** (unlike **NP**). Nevertheless, Kabanets and Impagliazzo [KI] showed that finding a deterministic polynomial-time algorithm for POLYNOMIAL IDENTITY TESTING implies that either **NEXP** $\not\subseteq$ **P/poly** or that the PERMANENT does not have polynomial-sized arithmetic circuits, both of which are long-standing open problems in complexity theory. (See [AvM] for a simpler and somewhat stronger proof.) POLYNOMIAL IDENTITY TESTING is in **co-RP**, by Theorem 2.12, but is not known to be complete for any randomized complexity class. Of course, it is also of interest to find additional complete problems for **prBPP**:

Open Problem 8.10. Find combinatorial or algebraic complete problems for any randomized complexity classes (e.g. **prBPP**, **prRP**, **prAM**, **BPP**, **RP**, **ZPP**).

Derandomizations of **prAM** are also known to imply circuit lower bounds, which are stronger than what the aforementioned results give from derandomizations of **prBPP** in that they either yield exponential-size bounds [AGHK] or give lower bounds for nondeterministic circuits [AvM].

One interesting interpretation of many of these results is to show that if derandomization is possible via any means (for all of **prBPP** or **prAM**), then it can be done in a canonical way — via pseudorandom generators (because these results show that derandomization implies circuit lower bounds, which in turn imply pseudorandom generators via Theorem 7.63). A recent work by Goldreich [Gol9] directly proves equivalences between various kinds of derandomizations of **prBPP** (e.g. worst-case or average-case) and various forms of pseudorandom generators, without going through circuit lower bounds. (See Problem 7.10.)

Another reason for the interest in these results is that they suggest derandomization as a potential approach to proving circuit lower bounds. Indeed, derandomization results have played a role in some state-of-the-art lower bounds, namely the result of Buhrman, Fortnow, and Thierauf [BFT] that **MAEXP** (the exponential-time analogue of **MA**) is not in **P/poly**, and the result of Williams [Wil1, Wil2] that **NEXP** is not in **ACC** (which is defined like **AC⁰** but also allowing unbounded fan-in gates that test whether their inputs sum to zero modulo m , for any constant m). The result of Kabanets and Impagliazzo [KI] (along with [AV]) has also been one of the motivations for the recent line of work on derandomizing POLYNOMIAL IDENTITY TESTING for low-depth arithmetic circuits. (See the Notes and References of Chapter 2.)

8.2.3 Hardness Amplification

As discussed in Section 7.6.1, *hardness amplification* is the task of taking a computational problem that is mildly hard on average and making it much more hard on average. Hardness amplification was introduced by Yao, in oral presentations of his paper [Yao]. Specifically, he suggested the “Direct Product” construction $f'(x_1, \dots, x_k) = (f(x_1), \dots, f(x_k))$ to convert a weak one-way function f (one that is mildly hard to invert) into a standard “strong” one-way function f' (satisfying Definition 7.10), and an “XOR Lemma” showing that if a Boolean function f is mildly hard to compute, then $f'(x_1, \dots, x_k) = f(x_1) \oplus f(x_2) \oplus \dots \oplus f(x_k)$ is very hard on average to compute. These were tools used in his proof that weak one-way permutations imply pseudorandom generators. These results have been generalized and strengthened in a number of ways:

Quantitative Bounds: It is of interest to have tight bounds on the hardness of f' as a function of the hardness of f and k . For the Direct Product construction, if f is δ average-case-hard, then intuitively we expect f' to be roughly $(1 - (1 - \delta)^k)$ average-case-hard, corresponding to the fact that an efficient algorithm trying to compute f' should have probability at most $1 - \delta$ of solving each of the k instances correctly. Similarly, in the case of the XOR Lemma, we expect that if f is $(1 - \delta)/2$ average-case-hard, then f' should be roughly $(1 - \delta^k)/2$ -average-case hard. Levin [Lev] proved a version of the XOR Lemma that gives such a tight bound, and his proof technique also extends to the Direct Product construction (see [GNW]).

Derandomization: Goldreich, Impagliazzo, Levin, Venkatesan, and Zuckerman [GIL⁺] gave derandomized hardness amplification results for converting weak one-way permutations and weak regular one-way functions into strong ones, where the inputs x_i are not independent but are generated in some pseudorandom way from a short seed that is the input to f' . Impagliazzo and Wigderson [Imp1, IW2] gave derandomized versions of the XOR Lemma and a Direct Product Lemma (for hard-to-compute boolean functions). These results were used in the first proof, due to [IW2], that $\mathbf{P} = \mathbf{BPP}$ if \mathbf{E} requires exponential-size circuits. Recall that the proof we saw in Chapter 7 avoids hardness amplification, and instead goes directly from worst-case hardness to strong average-case hardness via locally list-decodable codes (following [STV]). Nevertheless, hardness amplification is still of interest because it can be implemented in lower complexity than worst-case-to-average-case amplification. Indeed, hardness amplification (starting from mild average-case hardness) can be implemented in polynomial time with oracle access to f , whereas Viola [Vio1] has shown that black-box worst-case-to-average-case amplification (per Definition 7.74) cannot be implemented in the polynomial-time hierarchy (due to needing to compute a list-decodable encoding of the truth table of f). Indeed, another line of work has investigated hardness amplification for functions in \mathbf{NP} .

Hardness Amplification in \mathbf{NP} : The study of this topic was initiated in the work of O'Donnell [O'D1]. The goal is to show that if \mathbf{NP} has a function f that is mildly hard on average, then it has a function f' that is very hard on average. Yao's XOR Lemma does not prove this because f being in \mathbf{NP} does not imply that $f'(x_1, \dots, x_k) = f(x_1) \oplus \dots \oplus f(x_k)$ is also in \mathbf{NP} , assuming that $\mathbf{NP} \neq \mathbf{co-NP}$ (which is commonly conjectured). Thus, O'Donnell examines constructions of the form $f'(x_1, \dots, x_k) = C(f(x_1), \dots, f(x_k))$ where C is an efficiently computable and *monotone* combining function (i.e. changing an input of C from 0 to 1 cannot change the output from 1 to 0). He characterizes the amplification properties of C in terms of its “noise stability,” thereby connecting the study of hardness amplification with the analysis of boolean functions (see [O'D2] for more on this topic). He uses this connection to find monotone functions C with nearly optimal amplification properties, namely ones that will ensure that the function f' is roughly $(1/2 - O(1/\sqrt{k}))$ -hard if it is obtained by combining k evaluations of f . Contrast this bound with the XOR Lemma, where C is the (non-monotone) parity function and ensures that f' is $(1/2 - 1/2^{\Omega(k)})$ -hard. Healy, Vadhan, and Viola [HVV] showed how to derandomize O'Donnell's construction, so that the inputs x_1, \dots, x_k can be generated in a correlated way by a much shorter input to f' . This allows for taking k to be exponential in the input length of f' , and for certain combining functions C , the function f' is still in \mathbf{NP} (using the ability of a nondeterministic computation to compute exponential-sized ORs).

As a result, assuming that f is mildly hard for nonuniform algorithms running in time $2^{\Omega(n)}$ (the “high end”), they obtain $f' \in \mathbf{NP}$ that is $(1/2 - 1/2^{(n')^{1/2}})$ -hard where n' is the input length of f' . A quantitative improvement was given by [LTW1, GMR⁺], replacing $2^{(n')^{1/2}}$ with $2^{n'/\text{polylog}(n')}$, but it remains open to achieve the optimal bound of $2^{\Omega(n')}$.

Uniform Reductions: Another line of work has sought to give hardness amplification results for uniform algorithms, similarly to the work on derandomization from uniform assumptions described in Section 8.2.2. Like with cryptographic pseudorandom generators, most of the hardness amplification results in the cryptographic setting, such as Yao’s original hardness amplification for one-way functions [Yao], also apply to uniform algorithms. In the non-cryptographic setting, a difficulty is that black-box hardness amplification corresponds to error-correcting codes that can be decoded from very large distances, such as $1/2 - \varepsilon$ in the case of binary codes, and at these distances unique decoding is impossible, so one must turn to list decoding and use some nonuniform advice to select the correct decoding from the list. (See Definition 7.74 and the discussion after it. For amplification from mild average-case hardness rather than worst-case hardness, the coding-theoretic interpretation is that the decoding algorithm only needs to recover a string that is very close to the original message, rather than exactly equal to the message [Imp2, Tre2]; this also requires list decoding for natural settings of parameters.) However, unlike the case of pseudorandom generator constructions, here the number of candidates in the list can be relatively small (e.g. $\text{poly}(1/\varepsilon)$), so a reasonable goal for a uniform algorithm is to produce a list of possible decodings, as in our definition of locally list-decodable codes (Definition 7.54). As observed in [IW3, TV2, Tre2], if we are interested in amplifying hardness for functions in natural complexity classes such as \mathbf{NP} or \mathbf{E} , we can use (non-black-box) “checkability” properties of the initial function f to select a good decoding from the list, and thereby obtain a fully uniform hardness amplification result.

Uniform local list-decoding algorithms for the Reed–Muller Code were given by [AS2, STV] (as covered in Chapter 7), and were used to give uniform worst-case-to-average-case hardness amplification results for \mathbf{E} and other complexity classes in [TV2]. Trevisan [Tre2, Tre4] initiated the study of uniform hardness amplification from mild average-case hardness, giving uniform analogues of some of the hardness amplification results from [Imp1, O’D1]. Impagliazzo, Jaiswal, Kabanets, and Wigderson [IJK, IJKW] gave nearly optimal uniform Direct Product Theorems and XOR Lemmas. The existing uniform amplification results still do not quite match the nonuniform results in two respects. First, the derandomizations are not quite as strong; in particular, it is not known how to achieve an optimal “high end” result, converting a function on n -bit inputs that is mildly average-case hard against algorithms that run in time $2^{\Omega(n)}$ into a function on $O(n)$ -bit inputs that is $(1/2 - 1/2^{\Omega(n)})$ -hard against time $2^{\Omega(n)}$. (In the nonuniform setting, this was achieved by [IW2].) Second, for hardness amplification in \mathbf{NP} , the existing uniform results only amplify a function that is mildly hard against algorithms running in time t to ones that are $(1/2 - 1/(\log t)^{\Omega(1)})$ -hard, rather than $(1/2 - 1/t^{\Omega(1)})$ -hard, which is achieved in the nonuniform case by [O’D1, HVV]. See [BOKS] for a coding-theoretic approach to closing this gap.

Other Cryptographic Primitives: There has also been a large body of work on security amplification for other kinds of cryptographic primitives and interactive protocols (where the goal of the adversary is much more complex than just computing or inverting a function). Describing

this body of work is beyond the scope of this survey, so we simply refer the interested reader to [DIJK, Chu, Tes] and the references therein.

A key component of many hardness amplification results mentioned above is the Hardcore Theorem of Impagliazzo [Imp1] and variants. In its basic form, this theorem states that if a function f is mildly hard on average, then there is a *hardcore set* of inputs, of noticeable density, on which the function is very hard on average. Thus, intuitively, hardness amplification occurs when we evaluate a function many times because we are likely to hit the hardcore set. Since Impagliazzo’s original paper, there have been a number of papers giving quantitative improvements to the Hardcore Theorem [KS, Hol, BHK]. In addition to its applications in hardness amplification, the Hardcore Theorem has been shown to be closely related to “boosting” in machine learning [KS], “dense model theorems” that have applications in additive number theory and leakage-resilient cryptography [RTTV, DP, Tre7], computational analogues of entropy [STV, BSW, RTTV, CKLR, LLT, FOR, VZ], and “regularity lemmas” in the spirit of Szemerédi’s Regularity Lemma in graph theory [TTV]. One downside of using the Hardcore Theorem is that the (black-box) reductions in its proofs inherently require a lot of nonuniform advice [LTW2]. As shown by Holenstein [Hol], this nonuniformity can often be avoided in cryptographic settings, where one can efficiently sample random pairs $(x, f(x))$; his “uniform” analogue of the Hardcore Theorem and variants have played a key role in simplifying and improving the construction of cryptographic pseudorandom generators from one-way functions [Hol, HRV, VZ].

8.2.4 Deterministic Extractors

As mentioned in Section 6.5, the study of deterministic extractors has remained active even after the introduction of seeded extractors. One reason is that many applications (such as in cryptography, distributed computing, and Monte Carlo simulation) really require high-quality random bits, we often only have access to physical sources of randomness, and the trick of enumerating the seeds of a seeded extractor does not work in these contexts. Another is that deterministic extractors for certain classes of sources often have other applications of interest (even in contexts where we allow sources of truly random bits). For these reasons, after nearly a decade of work on simulating **BPP** with weak sources and seeded extractors, there was a resurgence of interest in deterministic extractors for various classes of sources [CDH⁺, TV1].

One important class of sources are those that consist of a small number of independent k -sources, first studied by Chor and Goldreich [CG1] (following [SV, Vaz3], who gave extractors for several independent unpredictable-bit sources). In addition to their motivation for obtaining high-quality randomness, extractors for 2 independent sources are of interest because of connections to communication complexity and to Ramsey theory. (Textbooks on these topics are [KN] and [GRS], respectively, and their connections to extractors can be found in [Vaz3, CG1] and the references below.) In particular, a disperser for 2 independent k -sources of length n is *equivalent* to a *bipartite Ramsey graph* — a bipartite graph with N vertices on each side that contains no $K \times K$ bipartite clique or $K \times K$ bipartite independent set (for $N = 2^n$ and $K = 2^k$). Giving explicit constructions of Ramsey graphs that approach $K = O(\log N)$ bound given by the probabilistic method [Erd1] is a long-standing open problem posed by Erdős [Erd2].

Chor and Goldreich [CG1] gave extractors for 2 independent k -sources when $k > n/2$ (e.g. inner-product modulo 2), and there was no improvement in this bound for nearly 2 decades. Substantial

progress began again with the work of Barak, Impagliazzo, and Wigderson [BIW], who used new results in arithmetic combinatorics to construct extractors for a constant number of independent k -sources when $k = \delta n$ for an arbitrarily small constant $\delta > 0$. Specifically, they used the Sum-Product Theorem over finite fields [BKT], which says that for p prime and every subset $A \subset \mathbb{F}_p$ whose size is not too close to p , either the set $A + A$ of pairwise sums or the set $A \cdot A$ of pairwise products is of size significantly larger than $|A|$. Arithmetic (and additive) combinatorics have now been found to be closely connected to many topics in pseudorandomness and theoretical computer science more broadly; see the survey of Trevisan [Tre6] and the talks and lecture notes from the minicourse [BTW].

A series of subsequent works improved [BIW] using techniques based on seeded extractors and condensers, various forms of composition, and other new ideas, achieving in particular explicit extractors for 3 sources of min-entropy $k = \delta n$ [BKS⁺], extractors for a constant number of sources of min-entropy $k = n^\delta$ [Rao] (notably making no use of arithmetic combinatorics), and dispersers for 2 sources of min-entropy $k = n^{o(1)}$ [BRSW]. The latter result is a substantial improvement over the previous best explicit construction of Ramsey graphs [FW], which avoided cliques and independent sets of size $K = 2^{\sqrt{n}}$ and only applied to the nonbipartite case. Even with all of this progress, the following remains open:

Open Problem 8.11. For every constant $\delta > 0$, construct an explicit function $\text{Ext} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ such that for every two independent δn -sources X, Y , $\text{Ext}(X, Y)$ is ε -close to uniform on $\{0, 1\}$.

Another line of work considers classes of sources defined by some measure of “complexity,” for example capturing the complexity of generating a random sample of the source (from independent coin tosses). Explicit extractors have been constructed for a variety of models of space-bounded sources [Blu, Vaz2, TV1, KM1, KM2, KRVZ]. In particular, Kamp, Rao, Vadhan, and Zuckerman [KRVZ] show how to deterministically extract $\Omega(n)$ bits from any source of min-entropy $k = \Omega(n)$ generated by a (nonuniform) algorithm with $o(n)$ space; this results exploits connections to extracting from multiple independent sources (as discussed above) as well as from bit-fixing sources (see Section 6.5). Trevisan and Vadhan [TV1] suggest to consider sources generated by small Boolean circuits, and show that under strong complexity assumptions (similar to, but stronger than, the one in Theorem 7.68), there is a deterministic extractor computable in time $\text{poly}(s)$ that extracts randomness from sources of min-entropy $k = (1 - \Omega(1))n$ generated by circuits of size s , with error $\varepsilon = 1/s$. It is necessary that the extractor has higher computational complexity than the source,³ but the min-entropy and error bounds can potentially be much better:

Open Problem 8.12. Under plausible complexity assumptions, show that there exists a deterministic extractor computable in time $\text{poly}(s)$ that extracts randomness from sources of min-entropy k generated by circuits of size s , with error ε , for $k \leq n/2$ and/or $\varepsilon = s^{-\omega(1)}$. Alternatively, give evidence that no such extractor exists.

De and Watson [DW] and Viola [Vio4, Vio5] have recently obtained unconditional extractors for restricted classes of sampling circuits, such as \mathbf{NC}^0 (where each output bit depends on a constant

³Interestingly, for *condensers*, it no longer seems necessary that the extractor has higher complexity than the source [DRV].

number of input bits) [DW, Vio4] \mathbf{AC}^0 [Vio4], even for sublinear min-entropy. These results are based on a close connection between constructing extractors for a class of circuits and finding explicit distributions that are hard for circuits in the class to sample, the latter a topic that was also studied for the purpose of proving data structure lower bounds [Vio5].

It is also natural to look at sources that are low complexity in an algebraic sense. The simplest such model is that of *affine sources*, which are uniform over affine subspaces of \mathbb{F}^n for a finite field \mathbb{F} . If the subspace has dimension k , then the source is a flat source of min-entropy $k \log |\mathbb{F}|$. For the case of $\mathbb{F} = \mathbb{Z}_2$, there are now explicit extractors for affine sources of sublinear min-entropy $k = o(n)$ [Bou, Yeh, Li] and dispersers for affine sources of subpolynomial min-entropy $k = n^{o(1)}$ [Sha2]. For large fields \mathbb{F} (i.e. $|\mathbb{F}| = \text{poly}(n)$), there are extractors for subspaces of every dimension $k \geq 1$ [GR]. There have also been works on extractors for sources described by polynomials of degree larger than 1, either as the output distribution of a low-degree polynomial map [DGW] or the zero set of low-degree polynomial (i.e. an algebraic variety) [Dvi]

We remark that several of the state-of-art constructions for independent sources and affine sources, such as [BKS⁺, BRSW, Sha2], are quite complicated, using sophisticated compositions and/or machinery from arithmetic combinatorics. It is of interest to find simpler constructions; some progress has been made for affine sources in [BK, Li] and for independent sources in [BZ] (where the latter in fact gives a reductions from constructing 2-source extractors to constructing affine extractors).

8.2.5 Algebraic Pseudorandomness

Algebraic Measures of Pseudorandomness In this survey, we have mostly focused on statistical and computational measures of (pseudo)randomness, such as pairwise independence and computational indistinguishability, respectively. It is also of interest to consider more algebraic measures, because they can be often related to statistical and/or computational measures, may be convenient for analyzing algebraic constructions of pseudorandom objects, and can have applications of their own (e.g. in additive or arithmetic combinatorics).

One of the most basic and important algebraic measures of pseudorandomness is that of a *small-bias* space, introduced by Naor and Naor [NN]. Here a random variable $X = (X_1, \dots, X_n)$ taking values in $\{0, 1\}^n$ is said to be ε -biased iff for every nonempty subset $S \subset [n]$, we have $(1 - \varepsilon)/2 \leq \Pr[\bigoplus_{i \in S} X_i = 1] \leq (1 + \varepsilon)/2$. Naor and Naor [NN] presented an explicit generator $G : \{0, 1\}^d \rightarrow \{0, 1\}^n$ of seed length $d = O(\log(n/\varepsilon))$ such that $G(U_d)$ is ε -biased; thus G is a pseudorandom generator fooling all “parity tests.” (Simpler constructions with better constants are in [AGHP].) This generator has found a variety of applications, such as almost k -wise independent hash functions [NN, AGHP] (cf. Problem 3.4), pseudorandom generators for small-width branching programs [SZ2, SZ, GMR⁺], derandomization of specific algorithms [NN], and almost-linear-length probabilistically checkable proofs [BSSVW, BSGH⁺].

Small-bias generators have several equivalent formulations. When viewed appropriately, they are equivalent to *linear* error-correcting codes (as in Problem 5.4) over \mathbb{F}_2 in which every nonzero codeword has Hamming weight between $(1 - \varepsilon)/2$ and $(1 + \varepsilon)/2$ [NN, AGHP]. They are also equivalent to expanders with spectral expansion $1 - \varepsilon$ that have algebraic structure of a *Cayley graph* over the group $G = \mathbb{Z}_2^n$ [AR]. (In general, when G is a group and $S \subset G$, the Cayley graph is $|S|$ -regular digraph on vertex set G , where the neighbors of vertex g are $\{gs : s \in S\}$.) Finally, the small-

bias property is equivalent to X being a distribution on the group $G = \mathbb{Z}_2^n$ all of whose nontrivial Fourier coefficients are at most ε [NN]. Specifically, the fact that the $(1 - \varepsilon)/2 \leq \Pr[\bigoplus_{i \in S} X_i = 1] \leq (1 + \varepsilon)/2$ is equivalent to requiring that $|\mathbb{E}[\chi_S(X)]| \leq \varepsilon$, where $\chi_S(x) = (-1)^{\bigoplus_{i \in S} x_i}$ is the Fourier character of $G = \mathbb{Z}_2^n$ indexed by the set S . (For background on Fourier analysis over finite groups, see the book by Terras [Ter], and for a survey of applications in theoretical computer science, see [dW] and the references therein.) Thus, we see that even in an algebraic context, different types of pseudorandom objects (generators, codes, and expanders) are equivalent.

These different views of small-bias spaces suggest different generalizations. One is that of linear codes over larger finite fields, which have been studied extensively in coding theory and of which we've seen some examples (Problem 5.4). Another is to consider Cayley graphs and Fourier analysis over groups G other than \mathbb{Z}_2^n . Over abelian groups, this generalization is fairly direct; the Fourier coefficients of a small-bias space on a group G are exactly the eigenvalues of a corresponding Cayley graph over the group G . For many abelian groups, including those of the form $G = \mathbb{Z}_p^n$ for prime p , there are known explicit constructions of small-bias generators with seed length $O(\log \log |G| + \log(1/\varepsilon))$, corresponding to explicit Cayley expanders of spectral expansion $1 - \varepsilon$ and degree $\text{poly}(\log |G|, 1/\varepsilon)$ [Kat, AIK⁺, Eve, AMN, RSW, AM]. However, there are benefits in working with non-abelian groups, as Cayley expanders over abelian groups G require degree $\Omega(\log N)$, where $N = |G|$ is the number of vertices [AR]. The logarithmic degree lower bound for expanders over abelian groups also holds for the more general notion of *Schreier graphs*, where the group G acts as permutations on the vertex set V , and we connect a vertex $v \in V$ to $s(v)$ for every $s \in S$ for some subset $S \subset G$. On the other hand, many of the algebraic constructions of constant-degree expanders, including Ramanujan graphs and the others described in Section 4.3.1, are obtained as Cayley graphs or Schreier graphs over non-abelian groups. The spectral expansion of these constructions can be analyzed via group representation theory, which is more involved than Fourier analysis over abelian groups, because it deals with matrix-valued (rather than scalar-valued) functions. See the surveys by Hoory, Linal, and Wigderson [HLW] and Lubotzky [Lub2], the lecture notes of Tao [Tao], and the notes for Chapter 4 for more on this and other approaches to analyzing the expansion of Cayley and Schreier graphs.

In the above, we think of a small-bias generator (according to the original definition) as producing pseudorandom elements of the group $G = \mathbb{Z}_2^n$, and consider generalizations to different groups G . An alternative view is that the small-bias generator produces a sequence of bits, and the group $G = \mathbb{Z}_2^n$ only arises in defining what it means for the distribution on bit-strings to be pseudorandom. More generally, we can take $G = H^n$ for any finite group H , and consider a random variable $X = (X_1, \dots, X_n)$ taking values in $\{0, 1\}^n$ to be pseudorandom if for every $(h_1, \dots, h_n) \in H^n$, the distribution of $h_1^{X_1} h_2^{X_2} \dots h_n^{X_n}$ is statistically close to $h_1^{R_1} h_2^{R_2} \dots h_n^{R_n}$ where $R = (R_1, \dots, R_n)$ is uniformly distributed in $\{0, 1\}^n$. This notion is of interest in part because the function $f_h(x_1, \dots, x_n) = h_1^{x_1} \dots h_n^{x_n}$ can be computed by a read-once branching program of width $|H|$ (see Section 8.2.1). Thus constructing generators fooling such “group product programs” are a natural warm-up to constructing pseudorandom generators for space-bounded computation, and in fact are equivalent to constructing generators for permutation branching programs in the case of constant width [KNP].

Yet another type of generalization is obtained by expanding the class of distinguishers from linear functions (which over \mathbb{F}_2^n are simply parities) to higher-degree polynomials. A series of recent results has shown that the sum of d small-bias spaces (over \mathbb{F}^n for any finite field \mathbb{F}) cannot be distinguished

from uniform by polynomials of degree d [BV, Lov4, Vio3]. These results were inspired by and influenced work on the Gowers uniformity norms from arithmetic combinatorics [Gow1, Gow2], which can be viewed as providing a “higher-order Fourier analysis” and have found numerous applications in theoretical computer science. (See [Tre6, BTW].)

Explicit Constructions via Polynomial Evaluation. As we’ve seen in this survey, algebra also provides powerful tools for constructing pseudorandom objects whose definitions make no reference to algebra. One particularly useful paradigm in such constructions is polynomial evaluation. Specifically, we construct a pseudorandom object $\Gamma : \mathbb{F}^n \times \mathbb{E} \rightarrow \mathbb{F}^m$, where \mathbb{F} is a finite field and $\mathbb{E} \subset \mathbb{F}^t$ is a set of evaluation points, by setting $\Gamma(f, y) = (f_1(y), \dots, f_m(y))$, where we view $f \in \mathbb{F}^n$ as specifying a low-degree polynomial in t variables, from which we construct m related low-degree polynomials f_1, \dots, f_m that we evaluate at the seed y . Reed–Solomon and Reed–Muller Codes (Constructions 5.14 and 5.16) correspond to the case where $m = 1$ and we take $f_1 = f$ to be a univariate or multivariate polynomial, respectively.⁴ In Parvaresh–Vardy Codes (Construction 5.21), we took f to be a univariate polynomial and obtained the f_i ’s by powering f modulo an irreducible polynomial. In Folded Reed–Solomon Codes (Construction 5.23), $f_i(Y) = f(\gamma^{i-1}Y)$, so evaluating f_i at y amounts to evaluating f at the related point $\gamma^{i-1}y$. More generally, if our construction $\Gamma(f, y)$ is obtained by evaluating f at points $g_1(y), \dots, g_m(y)$ for linear (or low-degree) functions g_1, \dots, g_m , we can also view it as evaluating the polynomials $f_1 = f \circ g_1, \dots, f_m = f \circ g_m$ at the seed y .

Prior to Parvaresh–Vardy codes, constructions of this type were used for extractors and pseudorandom generators. Specifically, Miltersen and Vinodchandran [MV] show that if we take f to describe a multivariate polynomial (via low-degree extension) and evaluate it on the points of a random axis-parallel line through y , we obtain a hitting-set generator construction against nondeterministic circuits (assuming f has appropriate worst-case hardness for nondeterministic circuits); this construction has played a key role in derandomizations of **AM** under uniform assumptions [GST, SU3].⁵ Ta-Shma, Zuckerman, and Safra [TSZS] showed that a similar construction yields randomness extractors with seed length $(1 + o(1)) \log n$ for polynomially small min-entropy and polynomial entropy loss. Shaltiel and Umans [SU1, Uma] showed that evaluating the t -variate polynomial f at the points $y, \gamma y, \gamma^2 y, \dots, \gamma^{m-1} y$, where γ is a primitive element of the field of size $|\mathbb{F}|^t$ (which we associate with \mathbb{F}^t) yields both a very good extractor construction and an optimal construction of pseudorandom generators from worst-case hard functions. Notice that this construction is precisely a multivariate analogue of Folded Reed–Solomon Codes (which came afterwards [GR]). Recently, Kopparty, Saraf, and Yekhanin [KSY] have introduced yet another useful way of obtaining the “related polynomials” f_1, \dots, f_m , namely by taking derivatives of f ; this has yielded the first codes with rate approaching 1 while being locally (list-)decodable in sublinear time [KSY, Kop] as well as codes matching the optimal rate-distance tradeoff of Folded Reed–Solomon Codes [GW, Kop]. All of this suggests that polynomial evaluation may be a promising approach to obtaining a unified and near-optimal construction of pseudorandom objects (Open

⁴In our presentation of Reed–Solomon and Reed–Muller codes, we took $\mathbb{E} = \mathbb{F}^t$, but many of the properties of these codes also hold for appropriately chosen subsets of evaluation points.

⁵The constructions described here have some additional components in addition to the basic polynomial evaluation framework $\Gamma(f, y) = (f_1(y), \dots, f_m(y))$, for example the seed should also specify the axis along which the line is parallel in [MV] and a position in an “inner encoding” in [TSZS, SU1, Uma]. We ignore these components in this informal discussion.

Problem 8.5).

References

- [AvM] S. Aaronson and D. van Melkebeek. On circuit lower bounds from derandomization. *Theory of Computing. An Open Access Journal*, 7:177–184, 2011.
- [AFK] M. Abadi, J. Feigenbaum, and J. Kilian. On hiding information from an oracle. *Journal of Computer and System Sciences*, 39(1):21–50, 1989.
- [Adl] L. Adleman. Two theorems on random polynomial time. In *19th Annual Symposium on Foundations of Computer Science (Ann Arbor, Mich., 1978)*, pages 75–83. IEEE, Long Beach, Calif., 1978.
- [Agr] M. Agrawal. On Derandomizing Tests for Certain Polynomial Identities. In *IEEE Conference on Computational Complexity*, pages 355–. IEEE Computer Society, 2003.
- [AB] M. Agrawal and S. Biswas. Primality and identity testing via Chinese remaindering. *Journal of the ACM*, 50(4):429–443 (electronic), 2003.
- [AKS] M. Agrawal, N. Kayal, and N. Saxena. PRIMES is in P. *Annals of Mathematics. Second Series*, 160(2):781–793, 2004.
- [AV] M. Agrawal and V. Vinay. Arithmetic Circuits: A Chasm at Depth Four. In *FOCS*, pages 67–75. IEEE Computer Society, 2008.
- [AIK⁺] M. Ajtai, H. Iwaniec, J. Komlós, J. Pintz, and E. Szemerédi. Construction of a thin set with small Fourier coefficients. *Bull. London Math. Soc.*, 22(6):583–590, 1990.
- [AKS1] M. Ajtai, J. Komlós, and E. Szemerédi. Sorting in $c \log n$ parallel steps. *Combinatorica*, 3(1):1–19, 1983.
- [AKS2] M. Ajtai, J. Komlós, and E. Szemerédi. Deterministic Simulation in LOGSPACE. In *19th Annual ACM Symposium on Theory of Computing*, pages 132–140, New York City, 25–27 May 1987.
- [AW] M. Ajtai and A. Wigderson. Deterministic Simulation of Probabilistic Constant Depth Circuits. In F. P. Preparata and S. Micali, editors, *Randomness and Computation*, volume 5 of *Advances in Computing Research*, pages 199–223. JAI Press Inc., 1989.
- [AKL⁺] R. Aleliunas, R. M. Karp, R. J. Lipton, L. Lovász, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *20th Annual Symposium on Foundations of Computer Science (San Juan, Puerto Rico, 1979)*, pages 218–223. IEEE, New York, 1979.
- [ABK⁺] E. Allender, H. Buhrman, M. Koucký, D. van Melkebeek, and D. Ronneburger. Power from random strings. *SIAM Journal on Computing*, 35(6):1467–1493 (electronic), 2006.
- [Alo1] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986. Theory of computing (Singer Island, Fla., 1984).
- [Alo2] N. Alon. Eigenvalues, geometric expanders, sorting in rounds, and Ramsey theory. *Combinatorica*, 6(3):207–219, 1986.
- [AC1] N. Alon and M. R. Capalbo. Explicit Unique-Neighbor Expanders. In *43rd Symposium on Foundations of Computer Science (Vancouver, BC, 2002)*, pages 73–79. IEEE, 2002.

- [AC2] N. Alon and F. R. K. Chung. Explicit construction of linear sized tolerant networks. *Discrete Mathematics*, 72(1-3):15–19, 1988. Proceedings of the First Japan Conference on Graph Theory and Applications (Hakone, 1986).
- [AGHP] N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple constructions of almost k -wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992. See also addendum in issue 4(1), 1993, pp. 199–210.
- [AGKS] N. Alon, V. Guruswami, T. Kaufman, and M. Sudan. Guessing secrets efficiently via list decoding. *ACM Transactions on Algorithms*, 3(4):Art. 42, 16, 2007.
- [AM] N. Alon and Y. Mansour. ϵ -discrepancy sets and their application for interpolation of sparse polynomials. *Inform. Process. Lett.*, 54(6):337–342, 1995.
- [AMS] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *Journal of Computer and System Sciences*, 58(1, part 2):137–147, 1999.
- [AM] N. Alon and V. D. Milman. Eigenvalues, Expanders and Superconcentrators (Extended Abstract). In *25th Annual Symposium on Foundations of Computer Science*, pages 320–322, Singer Island, Florida, 24–26 Oct. 1984. IEEE.
- [AR] N. Alon and Y. Roichman. Random Cayley graphs and expanders. *Random Structures & Algorithms*, 5(2):271–284, 1994.
- [AS1] N. Alon and J. H. Spencer. *The probabilistic method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience [John Wiley & Sons], New York, second edition, 2000. With an appendix on the life and work of Paul Erdős.
- [AS2] N. Alon and B. Sudakov. Bipartite subgraphs and the smallest eigenvalue. *Combinatorics, Probability and Computing*, 9(1):1–12, 2000.
- [ACR] A. E. Andreev, A. E. F. Clementi, and J. D. P. Rolim. Worst-Case Hardness Suffices for Derandomization: A New Method for Hardness-Randomness Trade-Offs. In P. Degano, R. Gorrieri, and A. Marchetti-Spaccamela, editors, *Automata, Languages and Programming, 24th International Colloquium*, volume 1256 of *Lecture Notes in Computer Science*, pages 177–187, Bologna, Italy, 7–11 July 1997. Springer-Verlag.
- [ACRT] A. E. Andreev, A. E. F. Clementi, J. D. P. Rolim, and L. Trevisan. Weak random sources, hitting sets, and BPP simulations. *SIAM Journal on Computing*, 28(6):2103–2116 (electronic), 1999.
- [AL] D. Angluin and D. Lichtenstein. Provable Security of Cryptosystems: A Survey. Technical Report YALEU/DCS/TR-288, Yale University, Department of Computer Science, 1983.
- [ALRS] S. Ar, R. J. Lipton, R. Rubinfeld, and M. Sudan. Reconstructing Algebraic Functions from Mixed Data. *SIAM Journal on Computing*, 28(2):487–510, 1999.
- [ASWZ] R. Armoni, M. Saks, A. Wigderson, and S. Zhou. Discrepancy sets and pseudorandom generators for combinatorial rectangles. In *37th Annual Symposium on Foundations of Computer Science (Burlington, VT, 1996)*, pages 412–421. IEEE Comput. Soc. Press, Los Alamitos, CA, 1996.
- [AB] S. Arora and B. Barak. *Computational complexity*. Cambridge University Press, Cambridge, 2009. A modern approach.
- [ALM⁺] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof Verification and the Hardness of Approximation Problems. *Journal of the ACM*, 45(3):501–555, May 1998.
- [AS1] S. Arora and S. Safra. Probabilistic Checking of Proofs: A New Characterization of NP. *Journal of the ACM*, 45(1):70–122, Jan. 1998.
- [AS2] S. Arora and M. Sudan. Improved Low Degree Testing and its Applications. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 485–495, El Paso, Texas, 4–6 May 1997.
- [Art] M. Artin. *Algebra*. Prentice Hall Inc., Englewood Cliffs, NJ, 1991.
- [AK] V. Arvind and J. Köbler. On pseudorandomness and resource-bounded measure. *Theoretical Computer Science*, 255(1-2):205–221, 2001.
- [AGHK] B. Aydinlioglu, D. Gutfreund, J. M. Hitchcock, and A. Kawachi. Derandomizing Arthur-Merlin games and approximate counting implies exponential-size lower bounds. *Computational Complexity*, 20(2):329–366, 2011.
- [AvM] B. Aydinlioglu and D. van Melkebeek. Nondeterministic Circuit Lower Bounds from Mildly Derandomizing Arthur-Merlin Games. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:80, 2012.
- [AMN] Y. Azar, R. Motwani, and J. Naor. Approximating probability distributions using small sample spaces. *Combinatorica*, 18(2):151–171, 1998.
- [BFLS] L. Babai, L. Fortnow, L. A. Levin, and M. Szegedy. Checking Computations in Polylogarithmic Time. In C. Koutsougeras and J. S. Vitter, editors, *STOC*, pages 21–31. ACM, 1991.
- [BFL] L. Babai, L. Fortnow, and C. Lund. Nondeterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1(1):3–40, 1991.

- [BFNW] L. Babai, L. Fortnow, N. Nisan, and A. Wigderson. BPP Has Subexponential Time Simulations Unless EXPTIME has Publishable Proofs. *Computational Complexity*, 3(4):307–318, 1993.
- [BM] L. Babai and S. Moran. Arthur-Merlin games: a randomized proof system, and a hierarchy of complexity classes. *Journal of Computer and System Sciences*, 36(2):254–276, 1988.
- [BHK] B. Barak, M. Hardt, and S. Kale. The uniform hardcore lemma via approximate Bregman projections. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1193–1200, Philadelphia, PA, 2009. SIAM.
- [BIW] B. Barak, R. Impagliazzo, and A. Wigderson. Extracting randomness using few independent sources. *SIAM Journal on Computing*, 36(4):1095–1118 (electronic), 2006.
- [BKS⁺] B. Barak, G. Kindler, R. Shaltiel, B. Sudakov, and A. Wigderson. Simulating independence: new constructions of condensers, Ramsey graphs, dispersers, and extractors. *Journal of the ACM*, 57(4):Art. 20, 52, 2010.
- [BRSW] B. Barak, A. Rao, R. Shaltiel, and A. Wigderson. 2-source dispersers for $n^{o(1)}$ entropy and Ramsey graphs beating the Frankl-Wilson construction. *Annals of Mathematics*, 2012. To appear. Preliminary version in *STOC '06*.
- [BSW] B. Barak, R. Shaltiel, and A. Wigderson. Computational analogues of entropy. In *Approximation, randomization, and combinatorial optimization*, volume 2764 of *Lecture Notes in Comput. Sci.*, pages 200–215. Springer, Berlin, 2003.
- [BTW] B. Barak, L. Trevisan, and A. Wigderson. Additive Combinatorics and Computer Science. <http://www.cs.princeton.edu/theory/index.php/Main/AdditiveCombinatoricsMinicourse>, August 2007.
- [BK] E. Barker and J. Kelsey. Recommendation for Random Number Generation Using Deterministic Random Bit Generators. National Institute of Standards and Technology, U.S. Department of Commerce, Special Publication 800-90A, January 2012.
- [Bas] L. A. Bassalygo. Asymptotically optimal switching circuits. *Problems of Information Transmission*, 17(3):206–211, 1981.
- [BSS] J. D. Batson, D. A. Spielman, and N. Srivastava. Twice-ramanujan sparsifiers. In *41st Annual ACM Symposium on Theory of Computing (Bethesda, MD)*, pages 255–262. ACM, 2009.
- [BF] D. Beaver and J. Feigenbaum. Hiding instances in multioracle queries (extended abstract). In *STACS 90 (Rouen, 1990)*, volume 415 of *Lecture Notes in Comput. Sci.*, pages 37–48. Springer, Berlin, 1990.
- [BGG] M. Bellare, O. Goldreich, and S. Goldwasser. Randomness in interactive proofs. *Computational Complexity*, 3(4):319–354, 1993.
- [BGM] M. Bellare, S. Goldwasser, and D. Micciancio. “Pseudo-Random” Number Generation Within Cryptographic Algorithms: The DDS Case. In B. S. K. Jr., editor, *CRYPTO*, volume 1294 of *Lecture Notes in Computer Science*, pages 277–291. Springer, 1997.
- [BR] M. Bellare and J. Rompel. Randomness-Efficient Oblivious Sampling. In *35th Annual Symposium on Foundations of Computer Science*, pages 276–287, Santa Fe, New Mexico, 20–22 Nov. 1994. IEEE.
- [BT] A. Ben-Aroya and A. Ta-Shma. A combinatorial construction of almost-Ramanujan graphs using the zig-zag product. In *40th Annual ACM Symposium on Theory of Computing (Victoria, British Columbia)*, pages 325–334. ACM, 2008.
- [BSGH⁺] E. Ben-Sasson, O. Goldreich, P. Harsha, M. Sudan, and S. Vadhan. Robust PCPs of proximity, shorter PCPs and applications to coding. *SIAM Journal on Computing*, 36(4):889–974, 2006.
- [BK] E. Ben-Sasson and S. Kopparty. Affine dispersers from subspace polynomials. In *STOC'09—Proceedings of the 2009 ACM International Symposium on Theory of Computing*, pages 65–74. ACM, New York, 2009.
- [BSSVW] E. Ben-Sasson, M. Sudan, S. Vadhan, and A. Wigderson. Randomness-efficient low degree tests and short PCPs via epsilon-biased sets. In *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, pages 612–621 (electronic), New York, 2003. ACM.
- [BZ] E. Ben-Sasson and N. Zewi. From affine to two-source extractors via approximate duality. In L. Fortnow and S. P. Vadhan, editors, *STOC*, pages 177–186. ACM, 2011.
- [BBR] C. H. Bennett, G. Brassard, and J.-M. Robert. Privacy amplification by public discussion. *SIAM Journal on Computing*, 17(2):210–229, 1988. Special issue on cryptography.
- [Ber1] S. J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Information Processing Letters*, 18(3):147–150, 1984.
- [Ber2] E. R. Berlekamp. *Algebraic coding theory*. McGraw-Hill Book Co., New York, 1968.
- [Ber3] E. R. Berlekamp. Factoring polynomials over large finite fields. *Mathematics of Computation*, 24:713–735, 1970.

- [BJKS] J. Bierbrauer, T. Johansson, G. Kabatianskii, and B. Smeets. On families of hash functions via geometric codes and concatenation. In *Advances in cryptology—CRYPTO '93 (Santa Barbara, CA, 1993)*, volume 773 of *Lecture Notes in Comput. Sci.*, pages 331–342. Springer, Berlin, 1994.
- [BL] Y. Bilu and N. Linial. Lifts, discrepancy and nearly optimal spectral gap. *Combinatorica*, 26(5):495–519, 2006.
- [Blu] M. Blum. Independent unbiased coin flips from a correlated biased source—a finite state Markov chain. *Combinatorica*, 6(2):97–108, 1986. Theory of computing (Singer Island, Fla., 1984).
- [BK] M. Blum and S. Kannan. Designing Programs that Check Their Work. *Journal of the ACM*, 42(1):269–291, 1995.
- [BLR] M. Blum, M. Luby, and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences*, 47(3):549–595, 1993.
- [BM] M. Blum and S. Micali. How to generate cryptographically strong sequences of pseudorandom bits. *SIAM Journal on Computing*, 13(4):850–864, 1984.
- [BDVY] A. Bogdanov, Z. Dvir, E. Verbin, and A. Yehudayoff. Pseudorandomness for Width 2 Branching Programs. *Electronic Colloquium on Computational Complexity (ECCC)*, 16:70, 2009.
- [BT] A. Bogdanov and L. Trevisan. Average-case complexity. *Foundations and Trends® in Theoretical Computer Science*, 2(1):1–106, 2006.
- [BV] A. Bogdanov and E. Viola. Pseudorandom bits for polynomials. *SIAM Journal on Computing*, 39(6):2464–2486, 2010.
- [BvzGH] A. Borodin, J. von zur Gathen, and J. Hopcroft. Fast parallel matrix and GCD computations. *Information and Control*, 52(3):241–256, 1982.
- [BD] C. Bosley and Y. Dodis. Does privacy require true randomness? In *Theory of cryptography*, volume 4392 of *Lecture Notes in Comput. Sci.*, pages 1–20. Springer, Berlin, 2007.
- [Bou] J. Bourgain. On the construction of affine extractors. *Geometric and Functional Analysis*, 17(1):33–57, 2007.
- [BKT] J. Bourgain, N. Katz, and T. Tao. A sum-product estimate in finite fields, and applications. *Geometric and Functional Analysis*, 14(1):27–57, 2004.
- [Boy] J. Boyar. Inferring sequences produced by pseudo-random number generators. *Journal of the Association for Computing Machinery*, 36(1):129–141, 1989.
- [Bra] M. Braverman. Polylogarithmic independence fools AC^0 circuits. *Journal of the ACM*, 57(5):Art. 28, 10, 2010.
- [BRRY] M. Braverman, A. Rao, R. Raz, and A. Yehudayoff. Pseudorandom Generators for Regular Branching Programs. In *FOCS*, pages 40–47. IEEE Computer Society, 2010.
- [Bro] A. Z. Broder. How hard is to marry at random? (On the approximation of the permanent). In *18th Annual ACM Symposium on Theory of Computing (Berkeley, CA)*, pages 50–58. ACM, 1986.
- [BV] J. Brody and E. Verbin. The Coin Problem and Pseudorandomness for Branching Programs. In *FOCS*, pages 30–39. IEEE Computer Society, 2010.
- [BF] H. Buhrman and L. Fortnow. One-sided two-sided error in probabilistic computation. In *STACS 99 (Trier)*, volume 1563 of *Lecture Notes in Comput. Sci.*, pages 100–109. Springer, Berlin, 1999.
- [BFT] H. Buhrman, L. Fortnow, and T. Thierauf. Nonrelativizing separations. In *Thirteenth Annual IEEE Conference on Computational Complexity (Buffalo, NY, 1998)*, pages 8–12. IEEE Computer Soc., Los Alamitos, CA, 1998.
- [BMRV] H. Buhrman, P. B. Miltersen, J. Radhakrishnan, and S. Venkatesh. Are bitvectors optimal? *SIAM Journal on Computing*, 31(6):1723–1744 (electronic), 2002.
- [BOKS] J. Buresh-Oppenheim, V. Kabanets, and R. Santhanam. Uniform Hardness Amplification in NP via Monotone Codes. *Electronic Colloquium on Computational Complexity (ECCC)*, 13(154), 2006.
- [BCS] P. Bürgisser, M. Clausen, and M. A. Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1997. With the collaboration of Thomas Lickteig.
- [CDH⁺] R. Canetti, Y. Dodis, S. Halevi, E. Kushilevitz, and A. Sahai. Exposure-Resilient Functions and All-Or-Nothing Transforms. In B. Preneel, editor, *Advances in Cryptology—EUROCRYPT 00*, Lecture Notes in Computer Science. Springer-Verlag, 14–18 May 2000.
- [CEG] R. Canetti, G. Even, and O. Goldreich. Lower bounds for sampling algorithms for estimating the average. *Information Processing Letters*, 53(1):17–25, 1995.
- [CRVW] M. Capalbo, O. Reingold, S. Vadhan, and A. Wigderson. Randomness Conductors and Constant-Degree Lossless Expanders. In *34th Annual ACM Symposium on Theory of Computing (STOC '02)*, pages 659–668, Montréal, CA, May 2002. ACM. Joint session with *CCC '02*.

- [CW] J. L. Carter and M. N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, 18(2):143–154, 1979.
- [Che1] J. Cheeger. A lower bound for the smallest eigenvalue of the Laplacian. In *Problems in analysis (Papers dedicated to Salomon Bochner, 1969)*, pages 195–199. Princeton Univ. Press, Princeton, N. J., 1970.
- [Che2] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23:493–507, 1952.
- [CG1] B. Chor and O. Goldreich. Unbiased Bits from Sources of Weak Randomness and Probabilistic Communication Complexity. *SIAM Journal on Computing*, 17(2):230–261, Apr. 1988.
- [CG2] B. Chor and O. Goldreich. On the power of two-point based sampling. *Journal of Complexity*, 5(1):96–106, 1989.
- [CGH⁺] B. Chor, O. Goldreich, J. Hästad, J. Friedman, S. Rudich, and R. Smolensky. The Bit Extraction Problem of t -Resilient Functions (Preliminary Version). In *FOCS*, pages 396–407. IEEE, 1985.
- [CGKS] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. *Journal of the ACM*, 45(6):965–982, 1998.
- [CG1] F. Chung and R. Graham. Sparse quasi-random graphs. *Combinatorica*, 22(2):217–244, 2002. Special issue: Paul Erdős and his mathematics.
- [CG2] F. Chung and R. Graham. Quasi-random graphs with given degree sequences. *Random Structures & Algorithms*, 32(1):1–19, 2008.
- [CGL] F. Chung, R. Graham, and T. Leighton. Guessing secrets. *Electronic Journal of Combinatorics*, 8(1):Research Paper 13, 25 pp. (electronic), 2001.
- [Chu] F. R. K. Chung. Diameters and eigenvalues. *Journal of the American Mathematical Society*, 2(2):187–196, 1989.
- [CGW] F. R. K. Chung, R. L. Graham, and R. M. Wilson. Quasi-random graphs. *Combinatorica*, 9(4):345–362, 1989.
- [Chu] K.-M. Chung. *Efficient Parallel Repetition Theorems with Applications to Security Amplification*. PhD thesis, Harvard University, 2011.
- [CKLR] K.-M. Chung, Y. T. Kalai, F.-H. Liu, and R. Raz. Memory delegation. In *Advances in cryptology—CRYPTO 2011*, volume 6841 of *Lecture Notes in Comput. Sci.*, pages 151–168. Springer, Heidelberg, 2011.
- [CW1] A. Cohen and A. Wigderson. Dispersers, Deterministic Amplification, and Weak Random Sources (Extended Abstract). In *30th Annual Symposium on Foundations of Computer Science (Research Triangle Park, North Carolina)*, pages 14–19. IEEE, 1989.
- [CW2] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *Journal of Symbolic Computation*, 9(3):251–280, 1990.
- [CLRS] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to algorithms*. MIT Press, Cambridge, MA, second edition, 2001.
- [CT] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley Series in Telecommunications. John Wiley & Sons, Inc., 2nd edition, 1991.
- [Csa] L. Csanky. Fast parallel matrix inversion algorithms. *SIAM Journal on Computing*, 5(4):618–623, 1976.
- [DSV] G. Davidoff, P. Sarnak, and A. Valette. *Elementary number theory, group theory, and Ramanujan graphs*, volume 55 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 2003.
- [De] A. De. Pseudorandomness for Permutation and Regular Branching Programs. In *IEEE Conference on Computational Complexity*, pages 221–231. IEEE Computer Society, 2011.
- [DW] A. De and T. Watson. Extractors and lower bounds for locally samplable sources. In *Approximation, randomization, and combinatorial optimization*, volume 6845 of *Lecture Notes in Comput. Sci.*, pages 483–494. Springer, Heidelberg, 2011.
- [dW] R. de Wolf. A Brief Introduction to Fourier Analysis on the Boolean Cube. *Theory of Computing, Graduate Surveys*, 1:1–20, 2008.
- [DL] R. A. DeMillo and R. J. Lipton. A Probabilistic Remark on Algebraic Program Testing. *Information Processing Letters*, 7(4):193–195, 1978.
- [DH] W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, 1976.
- [Din] I. Dinur. The PCP theorem by gap amplification. *Journal of the ACM*, 54(3):article 12, 44 pages (electronic), 2007.
- [DIJK] Y. Dodis, R. Impagliazzo, R. Jaiswal, and V. Kabanets. Security amplification for *interactive* cryptographic primitives. In *Theory of cryptography*, volume 5444 of *Lecture Notes in Comput. Sci.*, pages 128–145. Springer, Berlin, 2009.

- [DRV] Y. Dodis, T. Ristenpart, and S. Vadhan. Randomness Condensers for Efficiently Samplable, Seed-Dependent Sources. In R. Cramer, editor, *Proceedings of the 9th IACR Theory of Cryptography Conference (TCC '12)*, volume 7194 of *Lecture Notes on Computer Science*, pages 618–635. Springer-Verlag, 19–21 March 2012.
- [DP] D. Dubhashi and A. Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.
- [Dvi] Z. Dvir. Extractors for Varieties. In *IEEE Conference on Computational Complexity*, pages 102–113. IEEE Computer Society, 2009.
- [DGW] Z. Dvir, A. Gabizon, and A. Wigderson. Extractors and rank extractors for polynomial sources. *Computational Complexity*, 18(1):1–58, 2009.
- [DKSS] Z. Dvir, S. Kopparty, S. Saraf, and M. Sudan. Extensions to the method of multiplicities, with applications to Kakeya sets and mergers. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2009)*, pages 181–190. IEEE Computer Soc., Los Alamitos, CA, 2009.
- [DL] Z. Dvir and S. Lovett. Subspace evasive sets. In H. J. Karloff and T. Pitassi, editors, *STOC*, pages 351–358. ACM, 2012.
- [DS] Z. Dvir and A. Shpilka. Locally decodable codes with two queries and polynomial identity testing for depth 3 circuits. *SIAM Journal on Computing*, 36(5):1404–1434 (electronic), 2006/07.
- [DP] S. Dziembowski and K. Pietrzak. Leakage-Resilient Cryptography. In *FOCS*, pages 293–302. IEEE Computer Society, 2008.
- [Efr] K. Efremenko. 3-query locally decodable codes of subexponential length. In *STOC'09—Proceedings of the 2009 ACM International Symposium on Theory of Computing*, pages 39–44. ACM, New York, 2009.
- [Eli1] P. Elias. *List decoding for noisy channels*. Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, Mass., Rep. No. 335, 1957.
- [Eli2] P. Elias. The Efficient Construction of an Unbiased Random Sequence. *The Annals of Mathematical Statistics*, 43(3):865–870, June 1972.
- [Eli3] P. Elias. Error-correcting codes for list decoding. *IEEE Transactions on Information Theory*, 37(1):5–12, 1991.
- [Erd1] P. Erdős. Some remarks on the theory of graphs. *Bulletin of the American Mathematical Society*, 53:292–294, 1947.
- [Erd2] P. Erdős. Problems and results in chromatic graph theory. In *Proof Techniques in Graph Theory (Proc. Second Ann Arbor Graph Theory Conf., Ann Arbor, Mich., 1968)*, pages 27–35. Academic Press, New York, 1969.
- [EFF] P. Erdős, P. Frankl, and Z. Füredi. Families of finite sets in which no set is covered by the union of r others. *Israel Journal of Mathematics*, 51(1-2):79–89, 1985.
- [Eve] G. Even. Construction of Small Probabilistic Spaces for Deterministic Simulation. Master’s thesis, The Technion, 1991.
- [EGL⁺] G. Even, O. Goldreich, M. Luby, N. Nisan, and B. Velickovic. Efficient approximation of product distributions. *Random Struct. Algorithms*, 13(1):1–16, 1998.
- [ESY] S. Even, A. L. Selman, and Y. Yacobi. The complexity of promise problems with applications to public-key cryptography. *Information and Control*, 61(2):159–173, 1984.
- [FGL⁺] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. Interactive proofs and the hardness of approximating cliques. *Journal of the ACM*, 43(2):268–292, 1996.
- [Fil] J. A. Fill. Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, with an application to the exclusion process. *Annals of Applied Probability*, 1(1):62–87, 1991.
- [For] G. D. Forney. *Concatenated Codes*. MIT Press, 1966.
- [FW] P. Frankl and R. M. Wilson. Intersection theorems with geometric consequences. *Combinatorica*, 1(4):357–368, 1981.
- [FKS] M. L. Fredman, J. Komlós, and E. Szemerédi. Storing a sparse table with $O(1)$ worst case access time. *Journal of the ACM*, 31(3):538–544, 1984.
- [Fri] J. Friedman. A proof of Alon’s second eigenvalue conjecture and related problems. *Memoirs of the American Mathematical Society*, 195(910):viii+100, 2008.
- [FHK⁺] A. M. Frieze, J. Håstad, R. Kannan, J. C. Lagarias, and A. Shamir. Reconstructing truncated integer variables satisfying linear congruences. *SIAM Journal on Computing*, 17(2):262–280, 1988. Special issue on cryptography.
- [FOR] B. Fuller, A. O’Neill, and L. Reyzin. A Unified Approach to Deterministic Encryption: New Constructions and a Connection to Computational Entropy. In R. Cramer, editor, *TCC*, volume 7194 of *Lecture Notes in Computer Science*, pages 582–599. Springer, 2012.

- [GG] O. Gabber and Z. Galil. Explicit Constructions of Linear-Sized Superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, June 1981.
- [GR] A. Gabizon and R. Raz. Deterministic extractors for affine sources over large fields. *Combinatorica*, 28(4):415–440, 2008.
- [Gal] R. G. Gallager. *Low-Density Parity-Check Codes*. MIT Press, 1963.
- [GLR⁺] P. Gemmel, R. J. Lipton, R. Rubinfeld, M. Sudan, and A. Wigderson. Self-Testing/Correcting for Polynomials and for Approximate Functions. In C. Koutsougeras and J. S. Vitter, editors, *STOC*, pages 32–42. ACM, 1991.
- [GS] P. Gemmel and M. Sudan. Highly resilient correctors for polynomials. *Information Processing Letters*, 43(4):169–174, 1992.
- [Gil1] E. Gilbert. A comparison of signalling alphabets. *Bell Systems Technical Journal*, 31:504–522, 1952.
- [Gil2] J. Gill. Computational complexity of probabilistic Turing machines. *SIAM Journal on Computing*, 6(4):675–695, 1977.
- [Gil3] D. Gillman. A Chernoff bound for random walks on expander graphs. *SIAM Journal on Computing*, 27(4):1203–1220 (electronic), 1998.
- [GW] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995.
- [Gol1] O. Goldreich. A Sample of Samplers - A Computational Perspective on Sampling (survey). *Electronic Colloquium on Computational Complexity (ECCC)*, 4(20), 1997.
- [Gol2] O. Goldreich. *Modern cryptography, probabilistic proofs and pseudorandomness*, volume 17 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1999.
- [Gol3] O. Goldreich. *Foundations of cryptography*. Cambridge University Press, Cambridge, 2001. Basic tools.
- [Gol4] O. Goldreich. *Foundations of cryptography. II*. Cambridge University Press, Cambridge, 2004. Basic Applications.
- [Gol5] O. Goldreich. On promise problems: a survey. In *Theoretical computer science*, volume 3895 of *Lecture Notes in Comput. Sci.*, pages 254–290. Springer, Berlin, 2006.
- [Gol6] O. Goldreich. Probabilistic proof systems: a primer. *Foundations and Trends[®] in Theoretical Computer Science*, 3(1):1–91 (2008), 2007.
- [Gol7] O. Goldreich. *Computational complexity: a conceptual perspective*. Cambridge University Press, Cambridge, 2008.
- [Gol8] O. Goldreich. *A primer on pseudorandom generators*, volume 55 of *University Lecture Series*. American Mathematical Society, Providence, RI, 2010.
- [Gol9] O. Goldreich. In a World of P=BPP. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay of Randomness and Computation*, volume 6650 of *Lecture Notes in Computer Science*, pages 191–232. Springer, 2011.
- [GGM] O. Goldreich, S. Goldwasser, and S. Micali. How to Construct Random Functions. *Journal of the ACM*, 33(4):792–807, Oct. 1986.
- [GGR] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45(4):653–750, 1998.
- [GIL⁺] O. Goldreich, R. Impagliazzo, L. Levin, R. Venkatesan, and D. Zuckerman. Security preserving amplification of hardness. In *31st Annual Symposium on Foundations of Computer Science, Vol. I, II (St. Louis, MO, 1990)*, pages 318–326. IEEE Comput. Soc. Press, Los Alamitos, CA, 1990.
- [GKL] O. Goldreich, H. Krawczyk, and M. Luby. On the existence of pseudorandom generators. *SIAM Journal on Computing*, 22(6):1163–1175, 1993.
- [GL] O. Goldreich and L. A. Levin. A Hard-Core Predicate for all One-Way Functions. In *Proceedings of the Twenty First Annual ACM Symposium on Theory of Computing*, pages 25–32, Seattle, Washington, 15–17 May 1989.
- [GM] O. Goldreich and B. Meyer. Computational indistinguishability: algorithms vs. circuits. *Theoretical Computer Science*, 191(1-2):215–218, 1998.
- [GMW] O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity, or All languages in NP have zero-knowledge proof systems. *Journal of the ACM*, 38(3):691–729, 1991.
- [GNW] O. Goldreich, N. Nisan, and A. Wigderson. On Yao’s XOR lemma. Technical Report TR95–050, revision 2, Electronic Colloquium on Computational Complexity, June 2010. <http://www.eccc.uni-trier.de/eccc>.
- [GS] O. Goldreich and M. Sudan. Computational indistinguishability: a sample hierarchy. *Journal of Computer and System Sciences*, 59(2):253–269, 1999. 13th Annual IEEE Conference on Computation Complexity (Buffalo, NY, 1998).

- [GVW] O. Goldreich, S. Vadhan, and A. Wigderson. Simplified Derandomization of BPP Using a Hitting Set Generator. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay of Randomness and Computation*, volume 6650 of *Lecture Notes in Computer Science*, pages 59–67. Springer, 2011.
- [GW] O. Goldreich and A. Wigderson. Tiny Families of Functions with Random Properties: A Quality-Size Trade-off for Hashing. *Random Structures & Algorithms*, 11(4):315–343, 1997.
- [Gol] S. Goldwasser. Cryptography without (hardly any) secrets? In *Advances in cryptology—EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Comput. Sci.*, pages 369–370. Springer, Berlin, 2009.
- [GM] S. Goldwasser and S. Micali. Probabilistic Encryption. *Journal of Computer and System Sciences*, 28(2):270–299, Apr. 1984.
- [GMR1] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- [GMR2] P. Gopalan, R. Meka, and O. Reingold. DNF Sparsification and a Faster Deterministic Counting Algorithm. In *IEEE Conference on Computational Complexity*, pages 126–135. IEEE, 2012.
- [GMR⁺] P. Gopalan, R. Meka, O. Reingold, L. Trevisan, and S. Vadhan. Better Pseudorandom Generators via Milder Pseudorandom Restrictions. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS '12)*. IEEE, 20–23 October 2012. To appear.
- [Gow1] W. T. Gowers. A new proof of Szemerédi’s theorem for arithmetic progressions of length four. *Geometric and Functional Analysis*, 8(3):529–551, 1998.
- [Gow2] W. T. Gowers. A new proof of Szemerédi’s theorem. *Geometric and Functional Analysis*, 11(3):465–588, 2001.
- [GRS] R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey theory*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons Inc., New York, second edition, 1990. A Wiley-Interscience Publication.
- [Gur1] V. Guruswami. Guest column: error-correcting codes and expander graphs. *SIGACT News*, 35(3):25–41, 2004.
- [Gur2] V. Guruswami. *Algorithmic Results in List Decoding*, volume 2, number 2 of *Foundations and Trends in Theoretical Computer Science*. now publishers, 2006.
- [Gur3] V. Guruswami. Linear-Algebraic List Decoding of Folded Reed-Solomon Codes. In *IEEE Conference on Computational Complexity*, pages 77–85. IEEE Computer Society, 2011.
- [GHK] V. Guruswami, J. Håstad, and S. Kopparty. On the list-decodability of random linear codes. *IEEE Transactions on Information Theory*, 57(2):718–725, 2011.
- [GHSZ] V. Guruswami, J. Håstad, M. Sudan, and D. Zuckerman. Combinatorial bounds for list decoding. *IEEE Transactions on Information Theory*, 48(5):1021–1034, 2002.
- [GR] V. Guruswami and A. Rudra. Explicit codes achieving list decoding capacity: error-correction with optimal redundancy. *IEEE Transactions on Information Theory*, 54(1):135–150, 2008.
- [GS1] V. Guruswami and M. Sudan. Improved decoding of Reed-Solomon and algebraic-geometry codes. *Institute of Electrical and Electronics Engineers. Transactions on Information Theory*, 45(6):1757–1767, 1999.
- [GS2] V. Guruswami and M. Sudan. List decoding algorithms for certain concatenated codes. In *STOC*, pages 181–190, 2000.
- [GS3] V. Guruswami and M. Sudan. Extensions to the Johnson Bound. Unpublished Manuscript, February 2001.
- [GUV] V. Guruswami, C. Umans, and S. Vadhan. Unbalanced Expanders and Randomness Extractors from Parvaresh–Vardy Codes. *Journal of the ACM*, 56(4):1–34, 2009.
- [GW] V. Guruswami and C. Wang. Optimal rate list decoding via derivative codes. In *Approximation, randomization, and combinatorial optimization*, volume 6845 of *Lecture Notes in Comput. Sci.*, pages 593–604. Springer, Heidelberg, 2011.
- [GX] V. Guruswami and C. Xing. Folded codes from function field towers and improved optimal rate list decoding. In H. J. Karloff and T. Pitassi, editors, *STOC*, pages 339–350. ACM, 2012.
- [GST] D. Gutfreund, R. Shaltiel, and A. Ta-Shma. Uniform hardness versus randomness tradeoffs for Arthur-Merlin games. *Computational Complexity*, 12(3-4):85–130, 2003.
- [HRV] I. Haitner, O. Reingold, and S. Vadhan. Efficiency Improvements in Constructing Pseudorandom Generators from One-Way Functions. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC '10)*, pages 437–446, 6–8 June 2010.
- [Ham] R. W. Hamming. Error detecting and error correcting codes. *The Bell System Technical Journal*, 29:147–160, 1950.
- [HS] J. Hartmanis and R. E. Stearns. On the computational complexity of algorithms. *Transactions of the American Mathematical Society*, 117:285–306, 1965.

- [Har] N. J. A. Harvey. Algebraic Structures and Algorithms for Matching and Matroid Problems. In *47th Annual IEEE Symposium on Foundations of Computer Science (Berkeley, CA)*, pages 531–542. IEEE, 2006.
- [Hås] J. Håstad. *Computational Limitations of Small-Depth Circuits*. MIT Press, 1987.
- [HILL] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396 (electronic), 1999.
- [HVV] A. Healy, S. Vadhan, and E. Viola. Using nondeterminism to amplify hardness. *SIAM Journal on Computing*, 35(4):903–931 (electronic), 2006.
- [Hea] A. D. Healy. Randomness-efficient sampling within NC^1 . *Computational Complexity*, 17(1):3–37, 2008.
- [Hoe] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.
- [Hof] A. J. Hoffman. On eigenvalues and colorings of graphs. In *Graph Theory and its Applications (Proc. Advanced Sem., Math. Research Center, Univ. of Wisconsin, Madison, Wis., 1969)*, pages 79–91. Academic Press, New York, 1970.
- [Hol] T. Holenstein. Key agreement from weak bit agreement. In *STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 664–673. ACM, New York, 2005.
- [HLW] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the AMS*, 43(4):439–561, 2006.
- [Imp1] R. Impagliazzo. Hard-core distributions for somewhat hard problems. In *36th Annual Symposium on Foundations of Computer Science*, pages 538–545, Milwaukee, Wisconsin, 23–25 Oct. 1995. IEEE.
- [Imp2] R. Impagliazzo. Hardness as randomness: a survey of universal derandomization. In *Proceedings of the International Congress of Mathematicians, Vol. III (Beijing, 2002)*, pages 659–672, Beijing, 2002. Higher Ed. Press.
- [IJK] R. Impagliazzo, R. Jaiswal, and V. Kabanets. Approximate list-decoding of direct product codes and uniform hardness amplification. *SIAM Journal on Computing*, 39(2):564–605, 2009.
- [IJKW] R. Impagliazzo, R. Jaiswal, V. Kabanets, and A. Wigderson. Uniform direct product theorems: simplified, optimized, and derandomized. *SIAM Journal on Computing*, 39(4):1637–1665, 2009/10.
- [IKW] R. Impagliazzo, V. Kabanets, and A. Wigderson. In search of an easy witness: exponential time vs. probabilistic polynomial time. *Journal of Computer and System Sciences*, 65(4):672–694, 2002.
- [IR] R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *Advances in cryptology—CRYPTO '88 (Santa Barbara, CA, 1988)*, volume 403 of *Lecture Notes in Comput. Sci.*, pages 8–26. Springer, Berlin, 1990.
- [IW1] R. Impagliazzo and A. Wigderson. An information-theoretic variant of the inclusion- exclusion bound (preliminary version). Unpublished manuscript, 1996.
- [IW2] R. Impagliazzo and A. Wigderson. $P = BPP$ if E Requires Exponential Circuits: Derandomizing the XOR Lemma. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 220–229, El Paso, Texas, 4–6 May 1997.
- [IW3] R. Impagliazzo and A. Wigderson. Randomness vs time: derandomization under a uniform assumption. *Journal of Computer and System Sciences*, 63(4):672–688, 2001. Special issue on FOCS 98 (Palo Alto, CA).
- [IZ] R. Impagliazzo and D. Zuckerman. How to Recycle Random Bits. In *30th Annual Symposium on Foundations of Computer Science (Research Triangle Park, North Carolina)*, pages 248–253. IEEE, 1989.
- [IM] K. Iwama and H. Morizumi. An explicit lower bound of $5n - o(n)$ for Boolean circuits. In *Mathematical foundations of computer science 2002*, volume 2420 of *Lecture Notes in Comput. Sci.*, pages 353–364. Springer, Berlin, 2002.
- [JS] M. Jerrum and A. Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18(6):1149–1178, 1989.
- [JSV] M. Jerrum, A. Sinclair, and E. Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *Journal of the ACM*, 51(4):671–697 (electronic), 2004.
- [JM] S. Jimbo and A. Maruoka. Expanders obtained from affine transformations. *Combinatorica*, 7(4):343–355, 1987.
- [Jof1] A. Joffe. On a sequence of almost deterministic pairwise independent random variables. *Proceedings of the American Mathematical Society*, 29:381–382, 1971.
- [Jof2] A. Joffe. On a set of almost deterministic k -independent random variables. *Annals of Probability*, 2(1):161–162, 1974.
- [Joh1] S. Johnson. Upper bounds for constant weight error correcting codes. *Discrete Mathematics*, 3:109–124, 1972.

- [Joh2] S. M. Johnson. A new upper bound for error-correcting codes. *IRE Transactions on Information Theory*, IT-8:203–207, 1962.
- [Kab] V. Kabanets. Derandomization: A Brief Overview. In G. Paun, G. Rozenberg, and A. Salomaa, editors, *Current Trends in Theoretical Computer Science*, volume 1: Algorithms and Complexity, pages 165–188. World Scientific, 2004.
- [KI] V. Kabanets and R. Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1-2):1–46, 2004.
- [Kah] N. Kahale. Eigenvalues and expansion of regular graphs. *Journal of the ACM*, 42(5):1091–1106, 1995.
- [KLNS] J. D. Kahn, N. Linial, N. Nisan, and M. E. Saks. On the cover time of random walks on graphs. *Journal of Theoretical Probability*, 2(1):121–128, 1989.
- [Kal] A. T. Kalai. Unpublished manuscript, 2004.
- [KRVZ] J. Kamp, A. Rao, S. Vadhan, and D. Zuckerman. Deterministic extractors for small-space sources. *Journal of Computer and System Sciences*, 77(1):191–220, 2011.
- [KZ] J. Kamp and D. Zuckerman. Deterministic extractors for bit-fixing sources and exposure-resilient cryptography. *SIAM J. Comput.*, 36(5):1231–1247, 2006/07.
- [KPS] R. Karp, N. Pippenger, and M. Sipser. A time-randomness tradeoff. In *AMS Conference on Probabilistic Computational Complexity*, Durham, New Hampshire, 1985.
- [KL] R. M. Karp and R. J. Lipton. Turing machines that take advice. *L’Enseignement Mathématique. Revue Internationale. IIe Série*, 28(3-4):191–209, 1982.
- [KLM] R. M. Karp, M. Luby, and N. Madras. Monte Carlo approximation algorithms for enumeration problems. *Journal of Algorithms*, 10(3):429–448, 1989.
- [KUW] R. M. Karp, E. Upfal, and A. Wigderson. Constructing a perfect matching is in Random NC. *Combinatorica*, 6(1):35–48, 1986.
- [KL] J. Katz and Y. Lindell. *Introduction to modern cryptography*. Chapman & Hall/CRC Cryptography and Network Security. Chapman & Hall/CRC, Boca Raton, FL, 2008.
- [KT] J. Katz and L. Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 80–86 (electronic), New York, 2000. ACM.
- [Kat] N. M. Katz. An estimate for character sums. *Journal of the American Mathematical Society*, 2(2):197–200, 1989.
- [KS] N. Kayal and N. Saxena. Polynomial identity testing for depth 3 circuits. *Computational Complexity*, 16(2):115–138, 2007.
- [KvMS] J. Kinne, D. van Melkebeek, and R. Shaltiel. Pseudorandom generators, typically-correct derandomization, and circuit lower bounds. *Computational Complexity*, 21(1):3–61, 2012.
- [KS] A. R. Klivans and R. A. Servedio. Boosting and Hard-Core Set Construction. *Machine Learning*, 51(3):217–238, 2003.
- [KvM] A. R. Klivans and D. van Melkebeek. Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses. *SIAM Journal on Computing*, 31(5):1501–1526 (electronic), 2002.
- [Knu] D. E. Knuth. *The art of computer programming. Volume 2: Seminumerical Algorithms*. Addison–Wesley, third edition, 1998.
- [KM1] R. König and U. M. Maurer. Extracting randomness from generalized symbol-fixing and Markov sources. In *Proceedings of 2004 IEEE International Symposium on Information Theory*, page 232, 2004.
- [KM2] R. König and U. M. Maurer. Generalized Strong Extractors and Deterministic Privacy Amplification. In N. P. Smart, editor, *IMA Int. Conf.*, volume 3796 of *Lecture Notes in Computer Science*, pages 322–339. Springer, 2005.
- [Kop] S. Kopparty. List-Decoding Multiplicity Codes. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:44, 2012.
- [KSY] S. Kopparty, S. Saraf, and S. Yekhanin. High-rate codes with sublinear-time decoding. In L. Fortnow and S. P. Vadhan, editors, *STOC*, pages 167–176. ACM, 2011.
- [KNP] M. Koucký, P. Nimbhorkar, and P. Pudlák. Pseudorandom generators for group products: extended abstract. In L. Fortnow and S. P. Vadhan, editors, *STOC*, pages 263–272. ACM, 2011.
- [Kra] H. Krawczyk. How to predict congruential generators. *Journal of Algorithms*, 13(4):527–545, 1992.
- [KN] E. Kushilevitz and N. Nisan. *Communication complexity*. Cambridge University Press, Cambridge, 1997.
- [LR] O. Lachish and R. Raz. Explicit lower bound of $4.5n - o(n)$ for Boolean circuits. In *33rd Annual ACM Symposium on Theory of Computing*, pages 399–408 (electronic), New York, 2001. ACM.
- [Lan] H. O. Lancaster. Pairwise statistical independence. *Annals of Mathematical Statistics*, 36:1313–1317, 1965.
- [Lau] C. Lautemann. BPP and the polynomial hierarchy. *Information Processing Letters*, 17(4):215–217, 1983.

- [LLT] C.-J. Lee, C.-J. Lu, and S.-C. Tsai. Computational randomness from generalized hardcore sets. In *Fundamentals of computation theory*, volume 6914 of *Lecture Notes in Comput. Sci.*, pages 78–89. Springer, Heidelberg, 2011.
- [Lei] F. T. Leighton. *Introduction to parallel algorithms and architectures: arrays, trees, and hypercubes*. Morgan Kaufmann, San Mateo, CA, 1992.
- [Lev] L. A. Levin. One way functions and pseudorandom generators. *Combinatorica*, 7(4):357–363, 1987.
- [LV1] D. Lewin and S. Vadhan. Checking polynomial identities over any field: towards a derandomization? In *30th Annual ACM Symposium on the Theory of Computing (Dallas, TX)*, pages 438–447. ACM, New York, 1999.
- [LV2] M. Li and P. Vitányi. *An introduction to Kolmogorov complexity and its applications*. Texts in Computer Science. Springer, New York, third edition, 2008.
- [Li] X. Li. A New Approach to Affine Extractors and Dispersers. In *IEEE Conference on Computational Complexity*, pages 137–147. IEEE Computer Society, 2011.
- [LN] R. Lidl and H. Niederreiter. *Introduction to finite fields and their applications*. Cambridge University Press, Cambridge, first edition, 1994.
- [LLSZ] N. Linial, M. Luby, M. Saks, and D. Zuckerman. Efficient construction of a small hitting set for combinatorial rectangles in high dimension. *Combinatorica*, 17(2):215–234, 1997.
- [LN] N. Linial and N. Nisan. Approximate inclusion-exclusion. *Combinatorica*, 10(4):349–365, 1990.
- [Lip] R. J. Lipton. New directions in testing. In *Distributed computing and cryptography (Princeton, NJ, 1989)*, volume 2 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 191–202. Amer. Math. Soc., Providence, RI, 1991.
- [Lov1] L. Lovász. On determinants, matchings, and random algorithms. In *Fundamentals of Computation Theory (Berlin/Wendisch-Rietz)*, pages 565–574, 1979.
- [Lov2] L. Lovász. On the Shannon capacity of a graph. *IEEE Transactions on Information Theory*, 25(1):1–7, 1979.
- [Lov3] L. Lovász. *Combinatorial problems and exercises*. AMS Chelsea Publishing, Providence, RI, second edition, 2007.
- [Lov4] S. Lovett. Unconditional pseudorandom generators for low-degree polynomials. *Theory of Computing. An Open Access Journal*, 5:69–82, 2009.
- [Lu] C.-J. Lu. Improved pseudorandom generators for combinatorial rectangles. *Combinatorica*, 22(3):417–433, 2002.
- [LRVW] C.-J. Lu, O. Reingold, S. Vadhan, and A. Wigderson. Extractors: optimal up to constant factors. In *Proceedings of the 35th ACM Symposium on Theory of Computing (STOC '03)*, pages 602–611. ACM, 2003.
- [LTW1] C.-J. Lu, S.-C. Tsai, and H.-L. Wu. Improved hardness amplification in NP. *Theoretical Computer Science*, 370(1-3):293–298, 2007.
- [LTW2] C.-J. Lu, S.-C. Tsai, and H.-L. Wu. Complexity of hard-core set proofs. *Computational Complexity*, 20(1):145–171, 2011.
- [Lub1] A. Lubotzky. *Discrete groups, expanding graphs and invariant measures*, volume 125 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1994. With an appendix by Jonathan D. Rogawski.
- [Lub2] A. Lubotzky. Expander graphs in pure and applied mathematics. *American Mathematical Society. Bulletin. New Series*, 49(1):113–162, 2012.
- [LPS] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.
- [Lub1] M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15(4):1036–1053, 1986.
- [Lub2] M. Luby. Removing randomness in parallel computation without a processor penalty. *Journal of Computer and System Sciences*, 47(2):250–286, 1993.
- [LVW] M. Luby, B. Veličković, and A. Wigderson. Deterministic Approximate Counting of Depth-2 Circuits. In *ISTCS*, pages 18–24, 1993.
- [LW] M. Luby and A. Wigderson. *Pairwise Independence and Derandomization*, volume 1, number 4 of *Foundations and Trends in Theoretical Computer Science*. now publishers, 2005.
- [MS] F. J. MacWilliams and N. J. A. Sloane. *The theory of error-correcting codes*. North-Holland Publishing Co., Amsterdam, 1977. North-Holland Mathematical Library, Vol. 16.
- [Mar1] G. A. Margulis. Explicit constructions of expanders. *Problemy Peredači Informacii*, 9(4):71–80, 1973.
- [Mar2] G. A. Margulis. Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators. *Problemy Peredači Informacii*, 24(1):51–60, 1988.

- [MR] R. Martin and D. Randall. Disjoint decomposition of Markov chains and sampling circuits in Cayley graphs. *Combinatorics, Probability and Computing*, 15(3):411–448, 2006.
- [Mih] M. Mihail. Conductance and Convergence of Markov Chains-A Combinatorial Treatment of Expanders. In *30th Annual Symposium on Foundations of Computer Science (Research Triangle Park, North Carolina)*, pages 526–531. IEEE, 1989.
- [Mil1] G. L. Miller. Riemann’s Hypothesis and Tests for Primality. *Journal of Computer and System Sciences*, 13(3):300–317, Dec. 1976.
- [Mil2] P. Miltersen. *Handbook of Randomized Computing*, chapter Derandomizing Complexity Classes. Kluwer, 2001.
- [MV] P. B. Miltersen and N. V. Vinodchandran. Derandomizing Arthur-Merlin games using hitting sets. *Computational Complexity*, 14(3):256–279, 2005.
- [MU] M. Mitzenmacher and E. Upfal. *Probability and computing*. Cambridge University Press, Cambridge, 2005. Randomized algorithms and probabilistic analysis.
- [MR] R. Motwani and P. Raghavan. *Randomized algorithms*. Cambridge University Press, Cambridge, 1995.
- [MS] M. Mucha and P. Sankowski. Maximum Matchings via Gaussian Elimination. In *45th Symposium on Foundations of Computer Science (Rome, Italy)*, pages 248–255. IEEE Computer Society, 2004.
- [Mul] D. E. Muller. Boolean algebras in electric circuit design. *The American Mathematical Monthly*, 61(7, part II):27–28, 1954. Proceedings of the symposium on special topics in applied mathematics, Northwestern University (1953).
- [MVV] K. Mulmuley, U. V. Vazirani, and V. V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987.
- [Mut] S. Muthukrishnan. *Data Streams: Algorithms and Applications*, volume 1, number 2 of *Foundations and Trends in Theoretical Computer Science*. now publishers, 2005.
- [NN] J. Naor and M. Naor. Small-Bias Probability Spaces: Efficient Constructions and Applications. *SIAM Journal on Computing*, 22(4):838–856, Aug. 1993.
- [Nil] A. Nilli. On the second eigenvalue of a graph. *Discrete Mathematics*, 91(2):207–210, 1991.
- [Nis1] N. Nisan. Pseudorandom bits for constant depth circuits. *Combinatorica*, 11(1):63–70, 1991.
- [Nis2] N. Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992.
- [Nis3] N. Nisan. $RL \subseteq SC$. *Computational Complexity*, 4(1):1–11, 1994.
- [NT] N. Nisan and A. Ta-Shma. Extracting Randomness: A Survey and New Constructions. *Journal of Computer and System Sciences*, 58(1):148–173, February 1999.
- [NW] N. Nisan and A. Wigderson. Hardness vs Randomness. *Journal of Computer and System Sciences*, 49(2):149–167, Oct. 1994.
- [NZ] N. Nisan and D. Zuckerman. Randomness is Linear in Space. *Journal of Computer and System Sciences*, 52(1):43–52, Feb. 1996.
- [O’D1] R. O’Donnell. Hardness amplification within NP. *Journal of Computer and System Sciences*, 69(1):68–94, 2004.
- [O’D2] R. O’Donnell. Analysis of Boolean Functions. Book draft available at analysisofbooleanfunctions.org, 2012.
- [PV] F. Parvaresh and A. Vardy. Correcting Errors Beyond the Guruswami-Sudan Radius in Polynomial Time. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, pages 285–294. IEEE Computer Society, 2005.
- [Per] Y. Peres. Iterating von Neumann’s procedure for extracting random bits. *The Annals of Statistics*, 20(1):590–597, 1992.
- [Pet] W. W. Peterson. Encoding and error-correction procedures for the Bose-Chaudhuri codes. *IRE Transactions on Information Theory*, IT-6:459–470, 1960.
- [Pin] M. Pinsker. On the Complexity of a Concentrator. In *7th Annual Teletraffic Conference*, pages 318/1–318/4, Stockholm, 1973.
- [Pip] N. Pippenger. On Simultaneous Resource Bounds (Preliminary Version). In *20th Annual Symposium on Foundations of Computer Science (San Juan, Puerto Rico)*, pages 307–311. IEEE, 1979.
- [PF] N. Pippenger and M. J. Fischer. Relations among complexity measures. *Journal of the Association for Computing Machinery*, 26(2):361–381, 1979.
- [PB] R. L. Plackett and J. E. Burman. The design of optimum multi-factorial experiments. *Biometrika*, 33:305–325, 1945.
- [PHB] V. S. Pless, W. C. Huffman, and R. A. Brualdi, editors. *Handbook of coding theory. Vol. I, II*. North-Holland, Amsterdam, 1998.

- [Rab] M. O. Rabin. Probabilistic algorithm for testing primality. *Journal of Number Theory*, 12(1):128–138, 1980.
- [RT] J. Radhakrishnan and A. Ta-Shma. Bounds for dispersers, extractors, and depth-two superconcentrators. *SIAM Journal on Discrete Mathematics*, 13(1):2–24 (electronic), 2000.
- [Rag] P. Raghavan. Probabilistic construction of deterministic algorithms: approximating packing integer programs. *Journal of Computer and System Sciences*, 37(2):130–143, 1988. Twenty-Seventh Annual IEEE Symposium on the Foundations of Computer Science (Toronto, ON, 1986).
- [Ran] D. Randall. Mixing. In *44th Symposium on Foundations of Computer Science (Cambridge, MA)*, pages 4–15. IEEE Computer Society, 2003.
- [Rao] A. Rao. Extractors for a constant number of polynomially small min-entropy independent sources. *SIAM Journal on Computing*, 39(1):168–194, 2009.
- [RR] R. Raz and O. Reingold. On recycling the randomness of states in space bounded computation. In *Annual ACM Symposium on Theory of Computing (Atlanta, GA, 1999)*, pages 159–168 (electronic). ACM, New York, 1999.
- [RRV] R. Raz, O. Reingold, and S. Vadhan. Extracting all the Randomness and Reducing the Error in Trevisan’s Extractors. *Journal of Computer and System Sciences*, 65(1):97–128, August 2002.
- [RSW] A. Razborov, E. Szemerédi, and A. Wigderson. Constructing small sets that are uniform in arithmetic progressions. *Combin. Probab. Comput.*, 2(4):513–518, 1993.
- [Raz] A. A. Razborov. Lower bounds on the dimension of schemes of bounded depth in a complete basis containing the logical addition function. *Akademiya Nauk SSSR. Matematicheskie Zametki*, 41(4):598–607, 623, 1987.
- [RR] A. A. Razborov and S. Rudich. Natural proofs. *Journal of Computer and System Sciences*, 55(1, part 1):24–35, 1997. 26th Annual ACM Symposium on the Theory of Computing (STOC ’94) (Montreal, PQ, 1994).
- [Ree] I. S. Reed. A class of multiple-error-correcting codes and the decoding scheme. *IRE Transactions on Information Theory*, PGIT-4:38–49, 1954.
- [RS] I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. *Journal of the Society of Industrial and Applied Mathematics*, 8:300–304, 1960.
- [Rei1] O. Reingold. On black-box separations in cryptography. Tutorial at the Third Theory of Cryptography Conference (TCC ’06), March 2006. Slides available from <http://research.microsoft.com/en-us/people/omreing/>.
- [Rei2] O. Reingold. Undirected connectivity in log-space. *Journal of the ACM*, 55(4):Art. 17, 24, 2008.
- [RSW] O. Reingold, R. Shaltiel, and A. Wigderson. Extracting randomness via repeated condensing. *SIAM Journal on Computing*, 35(5):1185–1209 (electronic), 2006.
- [RTTV] O. Reingold, L. Trevisan, M. Tuliani, and S. Vadhan. Dense Subsets of Pseudorandom Sets. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS ’08)*, pages 76–85. IEEE, 26–28 October 2008.
- [RTV1] O. Reingold, L. Trevisan, and S. Vadhan. Notions of Reducibility between Cryptographic Primitives. In M. Naor, editor, *Proceedings of the First Theory of Cryptography Conference (TCC ’04)*, volume 2951 of *Lecture Notes in Computer Science*, pages 1–20. Springer-Verlag, 19–21 February 2004.
- [RTV2] O. Reingold, L. Trevisan, and S. Vadhan. Pseudorandom Walks in Regular Digraphs and the RL vs. L Problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC ’06)*, pages 457–466, 21–23 May 2006. Preliminary version as *ECCC TR05-22*, February 2005.
- [RVW1] O. Reingold, S. Vadhan, and A. Wigderson. Entropy Waves, the Zig-Zag Graph Product, and New Constant-Degree Expanders and Extractors. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS ’00)*, pages 3–13, Redondo Beach, CA, 17–19 Oct. 2000. IEEE.
- [RVW2] O. Reingold, S. Vadhan, and A. Wigderson. Entropy Waves, the Zig-Zag Graph Product, and New Constant-Degree Expanders. *Annals of Mathematics*, 155(1), January 2001.
- [RVW3] O. Reingold, S. Vadhan, and A. Wigderson. A Note on Extracting Randomness from Santha–Vazirani Sources. Unpublished manuscript, September 2004.
- [Rény] A. Rényi. On measures of entropy and information. In *Proc. 4th Berkeley Sympos. Math. Statist. and Prob., Vol. I*, pages 547–561. Univ. California Press, Berkeley, Calif., 1961.
- [RSA] R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the Association for Computing Machinery*, 21(2):120–126, 1978.
- [Ron] D. Ron. Property testing. In *Handbook of randomized computing, Vol. I, II*, volume 9 of *Comb. Optim.*, pages 597–649. Kluwer Acad. Publ., Dordrecht, 2001.

- [RV] E. Rozenman and S. Vadhan. Derandomized Squaring of Graphs. In *Proceedings of the 8th International Workshop on Randomization and Computation (RANDOM '05)*, number 3624 in Lecture Notes in Computer Science, pages 436–447, Berkeley, CA, August 2005. Springer.
- [Rub] R. Rubinfeld. Sublinear time algorithms. In *International Congress of Mathematicians. Vol. III*, pages 1095–1110. Eur. Math. Soc., Zürich, 2006.
- [RS] R. Rubinfeld and M. Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.
- [RSN⁺] A. Rukhin, J. Soto, J. Nechvatal, M. Smid, E. Barker, S. Leigh, M. Levenson, M. Vangel, D. Banks, A. Heckert, J. Dray, and S. Vo. A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications. National Institute of Standards and Technology, U.S. Department of Commerce, Special Publication 800-22, Revision 1a, April 2010.
- [SG] S. Sahni and T. Gonzalez. P -complete approximation problems. *Journal of the ACM*, 23(3):555–565, 1976.
- [SSZ] M. Saks, A. Srinivasan, and S. Zhou. Explicit OR-dispersers with polylogarithmic degree. *Journal of the ACM*, 45(1):123–154, 1998.
- [SZ1] M. Saks and S. Zhou. $\text{BP}_{\text{H}}\text{SPACE}(S) \subseteq \text{DSPACE}(S^{3/2})$. *Journal of Computer and System Sciences*, 58(2):376–403, 1999.
- [SZ2] M. Saks and D. Zuckerman. Unpublished manuscript, 1995.
- [SV] M. Sántha and U. V. Vazirani. Generating quasirandom sequences from semirandom sources. *Journal of Computer and System Sciences*, 33(1):75–87, 1986. Twenty-fifth annual symposium on foundations of computer science (Singer Island, Fla., 1984).
- [San] R. Santhanam. Circuit lower bounds for Merlin-Arthur classes. In *STOC'07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 275–283. ACM, New York, 2007.
- [Sar] P. Sarnak. *Some applications of modular forms*, volume 99 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1990.
- [Sav] W. J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4:177–192, 1970.
- [SSS] J. P. Schmidt, A. Siegel, and A. Srinivasan. Chernoff-Hoeffding bounds for applications with limited independence. *SIAM Journal on Discrete Mathematics*, 8(2):223–250, 1995.
- [Sch] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717, 1980.
- [Sha1] R. Shaltiel. Recent Developments in Extractors. In G. Paun, G. Rozenberg, and A. Salomaa, editors, *Current Trends in Theoretical Computer Science*, volume 1: Algorithms and Complexity, pages 189–228. World Scientific, 2004.
- [Sha2] R. Shaltiel. Dispersers for Affine Sources with Sub-polynomial Entropy. In R. Ostrovsky, editor, *FOCS*, pages 247–256. IEEE, 2011.
- [Sha3] R. Shaltiel. An introduction to randomness extractors. In *Automata, languages and programming. Part II*, volume 6756 of *Lecture Notes in Comput. Sci.*, pages 21–41. Springer, Heidelberg, 2011.
- [Sha4] R. Shaltiel. Weak derandomization of weak algorithms: explicit versions of Yao’s lemma. *Computational Complexity*, 20(1):87–143, 2011.
- [SU1] R. Shaltiel and C. Umans. Simple Extractors for All Min-Entropies and a New Pseudo-Random Generator. *Journal of the ACM*, 52(2):172–216, 2005.
- [SU2] R. Shaltiel and C. Umans. Pseudorandomness for approximate counting and sampling. *Computational Complexity*, 15(4):298–341, 2006.
- [SU3] R. Shaltiel and C. Umans. Low-end uniform hardness versus randomness tradeoffs for AM. *SIAM Journal on Computing*, 39(3):1006–1037, 2009.
- [Sha1] A. Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, 1979.
- [Sha2] A. Shamir. On the generation of cryptographically strong pseudorandom sequences. In *Automata, languages and programming (Akko, 1981)*, volume 115 of *Lecture Notes in Comput. Sci.*, pages 544–550. Springer, Berlin, 1981.
- [Sha3] C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 623–656, 1948.
- [SY] A. Shpilka and A. Yehudayoff. Arithmetic circuits: a survey of recent results and open questions. *Foundations and Trends[®] in Theoretical Computer Science*, 5(3-4):207–388 (2010), 2009.
- [SZ] J. Síma and S. Zák. Almost k -Wise Independent Sets Establish Hitting Sets for Width-3 1-Branching Programs. In A. S. Kulikov and N. K. Vereshchagin, editors, *CSR*, volume 6651 of *Lecture Notes in Computer Science*, pages 120–133. Springer, 2011.

- [Sip1] M. Sipser. A Complexity Theoretic Approach to Randomness. In *15th Annual ACM Symposium on Theory of Computing*, pages 330–335, Boston, Massachusetts, 25–27 Apr. 1983.
- [Sip2] M. Sipser. Expanders, randomness, or time versus space. *Journal of Computer and System Sciences*, 36(3):379–383, 1988. Structure in Complexity Theory Conference (Berkeley, CA, 1986).
- [Sip3] M. Sipser. *Introduction to the Theory of Computation*. Course Technology, 2nd edition, 2005.
- [SS] M. Sipser and D. A. Spielman. Expander codes. *IEEE Trans. Inform. Theory*, 42(6, part 1):1710–1722, 1996. Codes and complexity.
- [Smo] R. Smolensky. Algebraic Methods in the Theory of Lower Bounds for Boolean Circuit Complexity. In A. V. Aho, editor, *STOC*, pages 77–82. ACM, 1987.
- [SS] R. Solovay and V. Strassen. A fast Monte-Carlo test for primality. *SIAM Journal on Computing*, 6(1):84–85, 1977.
- [Spe] J. Spencer. *Ten lectures on the probabilistic method*, volume 64 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, second edition, 1994.
- [Spi] D. A. Spielman. Spectral Graph Theory and its Applications. In *48th Symposium on Foundations of Computer Science (FOCS 2007), 21-23 October 2007, Providence, RI, USA, Proceedings*, pages 29–38, 2007.
- [SZ] A. Srinivasan and D. Zuckerman. Computing with very weak random sources. *SIAM Journal on Computing*, 28(4):1433–1459 (electronic), 1999.
- [Ste1] T. Steinke. Pseudorandomness for Permutation Branching Programs Without the Group Theory. Technical Report TR12-083, Electronic Colloquium on Computational Complexity (ECCC), July 2012.
- [Ste2] J. Stern. Secret Linear Congruential Generators Are Not Cryptographically Secure. In *FOCS*, pages 421–426. IEEE Computer Society, 1987.
- [Sti1] H. Stichtenoth. *Algebraic function fields and codes*, volume 254 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, second edition, 2009.
- [Sti2] D. R. Stinson. Combinatorial techniques for universal hashing. *Journal of Computer and System Sciences*, 48(2):337–346, 1994.
- [Str] V. Strassen. Gaussian elimination is not optimal. *Numerische Mathematik*, 13:354–356, 1969.
- [Sud1] M. Sudan. Decoding of Reed Solomon Codes beyond the Error-Correction Bound. *Journal of Complexity*, 13(1):180–193, Mar. 1997.
- [Sud2] M. Sudan. Algorithmic Introduction to Coding Theory. Lecture notes, 2001. <http://people.csail.mit.edu/madhu/FT01/>.
- [Sud3] M. Sudan. Essential Coding Theory (Lecture Notes). <http://people.csail.mit.edu/madhu/FT04/>, 2004.
- [STV] M. Sudan, L. Trevisan, and S. Vadhan. Pseudorandom Generators without the XOR Lemma. *Journal of Computer and System Sciences*, 62:236–266, 2001.
- [TUZ] A. Ta-Shma, C. Umans, and D. Zuckerman. Lossless condensers, unbalanced expanders, and extractors. *Combinatorica*, 27(2):213–240, 2007.
- [TZ] A. Ta-Shma and D. Zuckerman. Extractor codes. *IEEE Transactions on Information Theory*, 50(12):3015–3025, 2004.
- [TSZS] A. Ta-Shma, D. Zuckerman, and S. Safra. Extractors from Reed-Muller codes. *Journal of Computer and System Sciences*, 72(5):786–812, 2006.
- [Tan] M. R. Tanner. Explicit Concentrators from Generalized N -gons. *SIAM Journal on Algebraic Discrete Methods*, 5(3):287–293, 1984.
- [Tao] T. Tao. Expansion in Finite Groups of Lie Type. Lecture Notes, 2012. <http://www.math.ucla.edu/~tao/254b.1.12w/>.
- [Ter] A. Terras. *Fourier analysis on finite groups and applications*, volume 43 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1999.
- [Tes] S. Tessaro. *Computational Indistinguishability Amplification*. PhD thesis, ETH Zurich, 2010. <http://e-collection.library.ethz.ch/eserv/eth:1817/eth-1817-02.pdf>.
- [Tre1] L. Trevisan. Extractors and pseudorandom generators. *Journal of the ACM*, 48(4):860–879 (electronic), 2001.
- [Tre2] L. Trevisan. List Decoding Using the XOR Lemma. In *Proceedings of the 44th IEEE Symposium on Foundations of Computer Science*, pages 126–135, Cambridge, MA, October 2003.
- [Tre3] L. Trevisan. Some Applications of Coding Theory in Computational Complexity. *Quaderni di Matematica*, 13:347–424, 2004.
- [Tre4] L. Trevisan. On uniform amplification of hardness in NP. In *STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 31–38. ACM, New York, 2005.

- [Tre5] L. Trevisan. Pseudorandomness and combinatorial constructions. In *International Congress of Mathematicians. Vol. III*, pages 1111–1136. Eur. Math. Soc., Zürich, 2006.
- [Tre6] L. Trevisan. Guest column: additive combinatorics and theoretical computer science. *SIGACT News*, 40(2):50–66, 2009.
- [Tre7] L. Trevisan. Dense model theorems and their applications. In *Theory of cryptography*, volume 6597 of *Lecture Notes in Comput. Sci.*, pages 55–57. Springer, Heidelberg, 2011.
- [TTV] L. Trevisan, M. Tulsiani, and S. Vadhan. Regularity, Boosting, and Efficiently Simulating Every High-Entropy Distribution. In *Proceedings of the 24th Annual IEEE Conference on Computational Complexity (CCC '09)*, pages 126–136, 15–18 July 2009. Preliminary version posted as *ECCC TR08-103*.
- [TV1] L. Trevisan and S. Vadhan. Extracting Randomness from Samplable Distributions. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS '00)*, pages 32–42, Redondo Beach, CA, 17–19 Oct. 2000. IEEE.
- [TV2] L. Trevisan and S. Vadhan. Pseudorandomness and Average-Case Complexity via Uniform Reductions. *Computational Complexity*, 16(4):331–364, December 2007.
- [Uma] C. Umans. Pseudo-random generators for all hardnesses. *Journal of Computer and System Sciences*, 67(2):419–440, 2003.
- [Vad] S. Vadhan. Probabilistic Proof Systems, Part I — Interactive & Zero-Knowledge Proofs. In S. Rudich and A. Wigderson, editors, *Computational Complexity Theory*, volume 10 of *IAS/Park City Mathematics Series*, pages 315–348. American Mathematical Society, 2004.
- [VZ] S. Vadhan and C. J. Zheng. Characterizing Pseudoentropy and Simplifying Pseudorandom Generator Constructions. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC '12)*, pages 817–836, 19–22 May 2012.
- [Vad] S. P. Vadhan. Constructing Locally Computable Extractors and Cryptosystems in the Bounded-Storage Model. *Journal of Cryptology*, 17(1):43–77, January 2004.
- [Val1] L. G. Valiant. Graph-theoretic properties in computational complexity. *Journal of Computer and System Sciences*, 13(3):278–285, 1976.
- [Val2] L. G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189–201, 1979.
- [Val3] L. G. Valiant. A Theory of the Learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- [Var] R. Varshamov. Estimate of the number of signals in error correcting codes. *Doklady Akad. Nauk SSSR*, 117:739–741, 1957.
- [Vaz1] U. V. Vazirani. Towards a Strong Communication Complexity Theory or Generating Quasi-Random Sequences from Two Communicating Slightly-random Sources (Extended Abstract). In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, pages 366–378, Providence, Rhode Island, 6–8 May 1985.
- [Vaz2] U. V. Vazirani. Efficiency Considerations in Using Semi-random Sources (Extended Abstract). In A. V. Aho, editor, *STOC*, pages 160–168. ACM, 1987.
- [Vaz3] U. V. Vazirani. Strong communication complexity or generating quasirandom sequences from two communicating semirandom sources. *Combinatorica*, 7(4):375–392, 1987.
- [VV] U. V. Vazirani and V. V. Vazirani. Random Polynomial Time Is Equal to Slightly-random Polynomial Time. In *26th Annual Symposium on Foundations of Computer Science*, pages 417–428, Portland, Oregon, 21–23 Oct. 1985. IEEE.
- [Vio1] E. Viola. The Complexity of Constructing Pseudorandom Generators from Hard Functions. *Computational Complexity*, 13(3-4):147–188, 2004.
- [Vio2] E. Viola. Pseudorandom bits for constant-depth circuits with few arbitrary symmetric gates. *SIAM Journal on Computing*, 36(5):1387–1403 (electronic), 2006/07.
- [Vio3] E. Viola. The sum of d small-bias generators fools polynomials of degree d . *Computational Complexity*, 18(2):209–217, 2009.
- [Vio4] E. Viola. Extractors for Circuit Sources. In R. Ostrovsky, editor, *FOCS*, pages 220–229. IEEE, 2011.
- [Vio5] E. Viola. The complexity of distributions. *SIAM Journal on Computing*, 41(1):191–218, 2012.
- [von] J. von Neumann. Various techniques used in conjunction with random digits. In *Collected works. Vol. V: Design of computers, theory of automata and numerical analysis*. The Macmillan Co., New York, 1963.
- [WC] M. N. Wegman and J. L. Carter. New hash functions and their use in authentication and set equality. *Journal of Computer and System Sciences*, 22(3):265–279, 1981.
- [Wil1] R. Williams. Improving exhaustive search implies superpolynomial lower bounds. In *STOC'10—Proceedings of the 2010 ACM International Symposium on Theory of Computing*, pages 231–240. ACM, New York, 2010.

- [Wil2] R. Williams. Non-uniform ACC Circuit Lower Bounds. In *IEEE Conference on Computational Complexity*, pages 115–125. IEEE Computer Society, 2011.
- [Woz] J. Wozencraft. List decoding. *Quarterly Progress Report, Research Laboratory of Electronics, MIT*, 48:90–95, 1958.
- [Yao] A. C. Yao. Theory and Applications of Trapdoor Functions (Extended Abstract). In *23rd Annual Symposium on Foundations of Computer Science*, pages 80–91, Chicago, Illinois, 3–5 Nov. 1982. IEEE.
- [Yeh] A. Yehudayoff. Affine extractors over prime fields. *Combinatorica*, 31(2):245–256, 2011.
- [Yek1] S. Yekhanin. Towards 3-query locally decodable codes of subexponential length. *Journal of the ACM*, 55(1):Art. 1, 16, 2008.
- [Yek2] S. Yekhanin. *Locally Decodable Codes*. Now Publishers, 2012. To appear.
- [Zip] R. Zippel. Probabilistic algorithms for sparse polynomials. In E. W. Ng, editor, *EUROSAM*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226. Springer, 1979.
- [Zuc1] D. Zuckerman. Simulating BPP Using a General Weak Random Source. *Algorithmica*, 16(4/5):367–391, Oct./Nov. 1996.
- [Zuc2] D. Zuckerman. Randomness-Optimal Oblivious Sampling. *Random Structures & Algorithms*, 11(4):345–367, 1997.
- [ZP] V. V. Zyablov and M. S. Pinsker. List cascade decoding (in Russian). *Problems of Information Transmission*, 17(4):29–33, 1981.