On the Hardness of Approximating

The Minimum Vertex Cover

and

The Closest Vector in a Lattice

Thesis submitted for the degree of "Doctor of Philosophy"
by
Irit Dinur

Submitted to the Senate of Tel-Aviv University
September 2001

The work on this thesis was carried out under the supervision of

Prof. Shmuel Safra

To my loving mother and father.

To Anat, my sister and masos.

To my dearest grandparents, Safta Galila,
Saba Menachem and Safta Runia,
and Saba Yehuda, whose love lives on.

To Hadar, who stood by me through some complex times, and to Chamy.

And to Yuvali, still in my heart.

Acknowledgment

I wish to thank my thesis advisor Prof. Muli Safra, for the endless hours and energy, for his guidance, and for his never-ending positivism.

Contents

Introduction

Computational Complexity Theory is the study of the complexity of computational problems, measured in terms of the amount of resources required for solving them.

A computational problem is a function, from the input to the output. Solving a computational problem means devising a systematic method, i.e. an algorithm, that finds the right output for every given input.

Optimization Problems

In an *optimization problem*, for each input there are many possible outputs, each assigned some numeric value by an *objective function.* An algorithm is said to solve an optimization problem if it finds, for any given input, an output of optimum value (minimum or maximum). Consider, for example, the Minimum-Vertex-Cover problem.

- *Input:* An undirected graph, $G = (V, E)$.
- *Output:* A subset $S \subseteq V$ touching all edges, that is, $\forall e \in E$, $e \cap S \neq \phi$.
- *Objective:* Minimize $|S|$.

The first question one asks, when faced with a computational problem, is whether or not there can be an efficient algorithm for it, i.e. is it *tractable.* As it turns out, many computational problems that arise naturally, including Vertex-Cover problem, are *NP-hard.* Thus, unless $P = NP$, there is no efficient algorithm for any of these problems.

Approximation Problems

For NP-hard optimization problems, the next best thing to finding the optimal solution would be to find an *approximate* solution that is, for example, at most twice the minimum (in case of a minimization problem). We say that an algorithm approximates the optimum to within a factor g, if it finds a solution that is provably no more than g times the optimum, in case of a minimization problem, or, in a maximization problem, at least the optimum divided by g.

For example, the following simple algorithm achieves a factor-2 approximation for the Minimum-Vertex-Cover problem. The algorithm constructs the set of vertices greedily, by adding at each step, both vertices of an edge that is not yet covered.

Great effort has been directed at finding approximation algorithms for problems that are NP-hard in their exact version. Some problems turn out to be approximable, while others are NP-hard even to approximate. Many times a problem yields itself to some loose approximation, yet remains NP-hard to approximate to within tighter (closer to 1) factors.

Gap Problems

The infeasibility of approximating a given optimization problem, is usually established by proving NP-hardness for the corresponding decision-like version of it, called the *gap* version. A gap-problem is a promise problem, with two parameters $g_1 < g_2$, such that the inputs are promised to be such that their optimal value is either greater than g_2 or smaller than g_1. An algorithm solving such a problem has to decide whether (assume a minimization problem) the minimum is smaller than g_1 (hence the instance is a 'yes' instance), or, whether the optimum is greater even than g_2 (and the instance is a 'no' instance). On instances whose minimum falls within the gap (g_2, g_1), the algorithm is allowed to return an arbitrary result; attributing the use of the term *gap*.

An algorithm approximating the optimum to within $\frac{g_1}{g_2}$ could distinguish 'yes' instances from 'no', according to whether or not its result exceeds g_2. Clearly, if the optimum is smaller than g_1, then the approximation algorithm could not have replied with a value larger than g_2. Thus, NP-hardness for the gap-version of a given problem, rules out the existence of an efficient approximation algorithm for it, unless $P = NP$.

PCP - Probabilistically Checkable Proofs

The complexity of many approximation problems is by now settled and a tight bound on the approximation ratio that can be efficiently achieved for them has been obtained. That is, it has been shown for these problems, that the approximation factor of the best known polynomial-time algorithm cannot be even marginally improved, unless P=NP. Hardness results for approximation problems, almost without exception, rely on the PCP characterization of NP as a starting point. The fundamental insight of the PCP characterization of NP is that it is NP-hard to distinguish between SAT formulas that are completely satisfied, and those that are *extremely* non-satisfiable, or, in other words, that *gap-SAT* is NP-hard. A SAT instance is a set of variables, and a set of *local-constraints* over them, each depending on the value of only a constant number of variables. The aim is to assign values to the variables that satisfy as many local-constraints as possible.

The PCP theorem was originally stated and proved in terms of proof systems. Imagine a verifier trying to verify correctness of some very long written proof. Rather than reading the entire proof, the verifier tosses some coins, and then probes the proof in a few places. Based on the values read from the proof, the verifier accepts or rejects the proof. The error probability of the verifier is the probability that the verifier accepts an incorrect proof, or rejects a correct one. [FGL$^+$91] discovered a connection between probabilistically checkable proofs, and approximation problems (specifically, Max-Clique). The PCP theorem [AS92, ALM$^+$92], quite surprisingly, showed that the class NP has such proofs, where the verifier reads only a constant number of bits and has constant error-probability. This was a breakthrough in the study of approximation problems, which brought about a flood of hardness of approximation results [ALM$^+$92, LY94, BGLR93, BGS98, Hås99, Hås97], to mention a few.

The formulation of the PCP theorem as the NP-hardness of gap-SAT is immediate, by taking one variable for each bit of the proof, and translating every test performed by the verifier to one local-constraint, imposed on variables corresponding to the verifier's probes, allowing only values that would make the verifier accept.

PCP characterizations of NP have served very well in resolving the complexity of approximation problems, leaving only a handful of classical optimization problems with the complexity of their approximation unsettled.

In this work we study the hardness of two combinatorial optimization problems, Minimum-Vertex-Cover and Closest-Vector, showing them NP-hard to approximate to within larger factors than previously known.

The Minimum Vertex Cover Problem

The Minimum Vertex Cover can be easily approximated to within a factor of 2, as described above. This can be only slightly improved, as the best known approximation algorithm [Hal00, BYE85, MS83] yields a factor only slightly smaller than 2. As to hardness results, the scheme of [BGS98, Hås99, Hås97], quite successful in achieving tight bounds for problems such as Max-3-Sat, Max-Linear-Equations, Max-Clique, was able to show Minimum-Vertex-Cover NP-hard to approximate to within a factor no larger than $\frac{7}{6}$, leaving open the gap between $\frac{7}{6}$ and 2.

The general scheme can be summarized as follows. Start with a gap-SAT instance, and replace each of the variables, with a set of variables representing its encoding, and each local-constraint with new constraints – whose form varies according to the problem at hand – that both verify the consistency of the encoding, and that the variables' encoded values satisfy the original local-constraints.

The encoding utilized in that scheme, as proposed in [BGS98], is the *long-code*, the most extensive binary code, whose bits correspond to all possible Boolean functions over the code's domain.

The p-biased–long-code. We introduce a generalization of the long-code, which we refer to as the *p-biased–long-code*, namely, a long-code on whose bits a probability distribution is superimposed. Considering non-uniform distributions, opens the door to new techniques for analyzing the long-code, clarifying some of its initially complicated structure. We describe a new connection between analyzing the long-code and the study of influences of variables on Boolean functions. This connection enables the use of known results from that field, in particular a very useful theorem of Friedgut [Fri98], for decoding the p-biased–long-code, extracting from it a few permissible values.

Composition. To utilize the p-biased–long-code, however, requires a redesign of the manner by which the Composition technique is applied. This new composition scheme starts with a phase in which the initial gap-SAT is preprocessed, coming up with a new set of symmetric variables whose consistency constraints are more suited for our purpose. An important feature of these variables, is that setting the value of one variable leaves at most two possible values for certain 'neighboring' variables. We then construct a graph, whose vertices correspond to encodings with the p-biased–long-code, of the assignments to these new variables.

Our proof must show that a small vertex-cover in this graph corresponds to a satisfying assignment to the initial gap-SAT instance. A difficulty arises from the fact that highlighting a small set of permissible values for our new variables, is not enough, and we must be able to distinguish one value for a significant portion of the variables, such that these distinguished values are consistent with each other.

This is achieved, with the aid of some lemmas from extremal set theory, that bound the size of *intersecting families*. This translates to show that if the minimum-vertex-cover in our graph is small enough, then it must distinguish *one* value for a significant portion of the variables. Based on these distinguished values we can derive global consistency, noting that as the initial consistency constraints between the new variables are rather loose, this only works provided $p < \frac{3-\sqrt{5}}{2} \approx 0.382$.

The Closest-Vector Problem

An n-dimensional lattice $L = L(v_1, .., v_n)$, for linearly independent vectors $v_1, .., v_n \in R^k$ is the additive group generated by the vectors, i.e. the set $L = \{\sum a_i v_i \mid a_i \in \mathbb{Z}\}$. Given L and an arbitrary vector y, the Closest Vector Problem (CVP) is to find a vector in L closest to y in a certain norm.

The best known polynomial-time algorithm approximating CVP, is an extension of the famous [LLL82] lattice-reduction due to [Sch85, Bab86], that approximates CVP to within $(1+\epsilon)^n$, for any constant $\epsilon > 0$. As for hardness of approximation results, CVP is known to be NP-hard to approximate to within any constant, and quasi-NP-hard to approximate to within $2^{\log^{(1-\epsilon)} n}$ for any constant $\epsilon > 0$ [ABSS93]. Still, there is quite a large gap between the smallest factor achievable by an algorithm and the largest factor for which hardness is known. Moreover, a third type of results [LLS90, GG98] provide indication that approximating these problems to within certain polynomial factors, is not NP-hard, unless the polynomial-time hierarchy collapses.

The proof of [ABSS93] utilizes amplification techniques, in which the dimension of the instance grows faster than the factor for which hardness of approximation is obtained. It is therefore unlikely that using this technique, even if allowing a super-polynomial blow-up, one can obtain hardness for factors larger than $2^{(\log n)^{1-\epsilon}}$ for any constant $\epsilon > 0$.

We improve on [ABSS93] in two ways. First, we go beyond the factor of $2^{(\log n)^{1-\epsilon}}$ for any constant $\epsilon > 0$, which was the previous hardness-of-approximation factor known for CVP. Instead, we achieve a factor of $2^{\frac{\log n}{\log \log n}} = n^{1/\log \log n}$. Furthermore, we show approximating CVP is *NP-hard* for these large factors, compared to the previously known *quasi* NP-hardness.

As the known PCP characterizations of NP seem inadequate for showing hardness of approximating CVP to within large factors, we introduce a new NP-hard gap problem, S-SAT. This problem is also a gap version of SAT, one with a different objective function. The main part of our proof is to establish the NP-hardness of S-SAT, and is carried out via similar techniques to the ones used for proving the PCP theorem. We apply a recursive composition, employing two alternate types of algebraic extensions, to achieve a gap of $n^{c/\log \log n}$.

This work is divided into two parts. Part I, devoted to the study of Minimum-Vertex-Cover, is based on [DS01]. Part II of this work, dealing with the Closest Vector Problem, is based on [DKS98, DKRS99].

Part I

Vertex Cover

Chapter 1

Introduction

A combinatorial optimization problem, that maybe most captures the limitations of current technique for proving hardness of approximation, is Minimum-Vertex-Cover: the problem of finding a smallest set of vertices that cover (touch) all edges in a given graph. The complement of a vertex cover must be an independent set, a set of vertices with no edges between them. Thus, finding the Minimum-Vertex-Cover is equivalent to finding the Maximum-Independent-Set. Where approximation is concerned, however, since the minimum vertex-cover may be much larger than the maximum independent set, a good approximation of the former does not carry on to the later. In fact, it has been proven to be hard to approximate the maximum independent set to within $|V|^{1-o(1)}$ [Hås99, EH00, Kho01], yet easy to approximate the minimum vertex cover to within a factor slightly smaller than 2 [Hal00, BYE85, MS83].

In what follows, we extend current technique, to show it is hard to approximate Minimum-Vertex-Cover to within a factor larger than the previously known $\frac{7}{6}$ [Hås97]. More specifically, a corollary of our analysis is the following:

Corollary 4.2 *Given a graph G, it is NP-hard to approximate the size of the smallest Vertex-Cover to within a factor of* 1.361. ■

Background. Let us now briefly describe the background related to PCP, hardness of approximation, and reductions utilizing one to obtain the other.

Minimum Vertex Cover belongs to the class APX-complete, defined in [PY91], of problems whose hardness of approximation is interrelated. The PCP theorem [AS92, ALM$^+$92] implies that it is NP-hard to approximate, to within *some* constant factor, every problem in this class. This, however, is far from providing a tight bound for these approximation problems, as the constant factor of approximation whose hardness is thus obtained is usually quite far from the known upper-bound. For tight bounds one is required to work a little harder, and sometimes devise ingenious reductions and elaborate analysis.

One of the most successful recipes for such reductions, is the scheme of [BGS98, Hås97, Hås99], whose rough sketch is as follows.

Given a gap-SAT instance Ψ, first apply the parallel repetition lemma of [Raz98]. This, for any parameter k, results in a new gap-SAT problem, which we refer to as $\mathsf{Par}\left[\Psi, k\right]$, over (non-binary) variables of two types X and Y, and in which each local-constraint depends on one variable of X and one of Y. $\mathsf{Par}\left[\Psi, k\right]$ has a satisfying assignment if Ψ has one, and, otherwise, not even an arbitrarily small (exponentially small in k) fraction of the constrains can be satisfied.

In the next step, one applies a version of the Composition technique of [AS92], as proposed in [BGS98], to this specific setting. The composition replaces each of the variables of X and Y, with a set of variables representing its encoding, and each local-constraint $\varphi(x, y) \in \mathsf{Par}\left[\Psi, k\right]$ with constraints that both verify the consistency of the encoding of x and y, and that x and y's encoded values satisfy $\varphi(x, y)$.

This new set of constraints may take different form according to the problem one intends to show hard. The next step of the reduction, if necessary, translates those local-constraints to an instance of the problem at hand, whose solution, even if only approximates the best solution, implies a satisfying assignment for $\mathsf{Par}\left[\Psi, k\right]$ and thereby for Ψ as well.

The encoding utilized in that scheme, as proposed in [BGS98], is the *long-code*, the most extensive binary code, whose bits correspond to all possible Boolean functions over the code's domain. Alternatively, the long-code can be represented as a sequence of subsets of R, specifying, for each bit, which of the elements of R has 1 on that bit of their encoding. The long-code is extremely inefficient in size, however, since the range of values variables in X and Y can take is rather small, this poses no problem. Numerous tight bounds for approximation problems, such as Max-3-Sat, Linear Equations, Max-Clique, have been thus obtained. Some of these involve an extensive analysis of consistency tests over long-codes, using Fourier analysis [Hås99, Hås97], showing it suffices, for example, to probe the value of only three bits of the long-codes of x and y to be assured, with high probability, of the consistency within the encoding of x and y as well as the consistency between the two.

Vertex Cover. Nevertheless, where other open questions regarding the hardness of approximation problems rise and fall, Vertex-Cover has stood still, leaving its best hardness result nowhere higher than the $\frac{7}{6}$ factor of [Hås97], which is still far from the best known upper bound [Hal00, BYE85, MS83] of a factor slightly smaller than 2.

Our analysis herein amends the [BGS98, Hås97] scheme in several ways, most notably by introducing a generalization of the long-code, which we refer to as the *biased-long-code*, namely, a long-code on whose bits a probability distribution is superimposed. The probability attached to each Boolean function (that is, a bit of the code), is determined by tossing a p-biased coin for each element in the domain, taking true with probability p, and false with probability $1 - p$.

Given an assignment A to the bits of a p-biased–long-code, the weight of A, namely, the fractional size of the set of bits assigned 1 by A, is determined according to this probability distribution. The original long-code is a special case of this construct, in which $p = \frac{1}{2}$ and the distribution over the bits is uniform.

To utilize the p-biased–long-code, however, requires a redesign of the manner by which the Composition technique is applied. This new scheme starts with a phase in which $\mathsf{Par}[\Psi, k]$ is preprocessed, coming up with a set of variables $\mathcal{Z}$ whose consistency constraints are quite loose, nevertheless, whose structure can still be utilized by our analysis to deduce global consistency assuming local-consistency, which, in this case, translates to a large independent set.

Overview of the Proof. Starting with a gap-SAT instance $\Phi = \mathsf{Par}[\Psi, k]$, we first preprocess it, coming up with a new set of *blocks* $\mathcal{Z}$ whose consistency constraints are more suited for our purpose. We then construct a graph $\mathcal{G}_{[p,p^\bullet]}(\Phi)$, whose vertices correspond to encodings with the p-biased–long-code, of the assignments to the blocks $\mathcal{Z}$. We next proceed to show that $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ has an independent set whose weight is $p-\varepsilon$ if Φ is satisfiable, or otherwise, that the weight of the largest independent set in $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ is at most $p^\bullet \overset{def}{=} \max(p^2, 4p^3 - 3p^4) + \varepsilon$ (provided $p < \frac{3-\sqrt{5}}{2} \approx 0.382$). While the first (completeness) part follows directly from the definition of the p-biased–long-code, the soundness part requires deeper analysis of assignments to the p-biased–long-code, and relies heavily on an extensive study of the *influence* of variables on Boolean functions. This study has been carried out for quite a while, in an impressive sequence of papers [BOL89, BOLS88, KKL88, BK97, FK96, BKS99], culminating in the result – which we make a good use of herein – of Friedgut [Fri98] (Theorem 2.3). Friedgut's lemma essentially asserts that Boolean functions of low average-sensitivity (namely Boolean functions that infrequently change value when one of their variables is flipped at random) are almost entirely determined by the values of only a small set of variables.

We are able to apply Friedgut's lemma and extract from the biased–long-code of each variable in $\mathcal{Z}$, a few permissible values. For this we must utilize additional properties of the graph $\mathcal{G}_{[p,p^\bullet]}(\Phi)$, to show that an independent set corresponds to an encoding with biased–long-code, whose average-sensitivity is low. An independent set in $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ identifies a few permissible values in most blocks, even without being larger than $p^\bullet + \varepsilon$. However, these values yield insufficiently weak consistency between the blocks, such that can be attained even when Φ is far from satisfiable. For true global consistency, we must venture into the field of extremal set theory to show that if the independent set in $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ is larger than $p^\bullet = \max(p^2, 4p^3 - 3p^4) + \varepsilon$, enough variables in $\mathcal{Z}$ each distinguish *one* value. Finally, based on these distinguished values, we can derive consistency between variables in $\mathcal{Z}$ noting that, as the initial consistency constraints between $\mathcal{Z}$'s variables are rather loose, this only works provided $p < \frac{3-\sqrt{5}}{2} \approx 0.382$.

Outline. Our presentation of the hardness result for Minimum Vertex Cover stretches across chapters 2-5. Chapter 2 is devoted to the various mathematical tools, used for analyzing the long-code, and for obtaining the necessary combinatorial upper bounds for deriving global consistency from local-consistency. Chapter 3 presents the reduction from PCP to Minimum Vertex Cover, namely, we show how to translate any gap-SAT instance Φ, to a graph $\mathcal{G}_{[p,p^\bullet]}(\Phi)$, such that approximating the Minimum Vertex Cover in $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ yields an approximation of the maximum fraction of satisfiable tests in Φ. Chapter 4 is devoted to the proof of the correctness of the result: the completeness and (mainly) the soundness of the reduction. We conclude this chapter showing that our analysis of $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ is tight, i.e. that an independent set of weight $p^\bullet$ exists even if Φ is not satisfiable.

We conclude with a short discussion of possible extensions of our technique, and applications to related open problems. In order to make the exposition of all these combinatorial

Chapter 2

The p-Biased Long Code

The long-code over a domain R encodes each element of R by the longest possible (without repetition) sequence of binary bits, corresponding to all possible Boolean functions over R. Each bit can be canonically identified with the subset of all elements of R whose encoding is 1 on that bit. Thus, the bits of the long-code become the power set of R, $\mathbf{P}(R) \overset{def}{=} \{F \subseteq R\}$.

Notation. As our analysis combines ideas from different fields, we denote, adopting notation from extremal set theory, a family of subsets of R by $\mathcal{F} \subseteq \mathbf{P}(R)$, and one subset in it by $F \in \mathcal{F}$.

Let us formally define the long-code of R,

Definition 2.1 *The* long-code *of R, denoted $\mathcal{LC}^R$, is the most extensive binary code, namely the code consisting of* all *subsets of R,*

$$\mathcal{LC}^R \overset{def}{=} \mathbf{P}(R) \ .$$

A codeword $E\colon \mathbf{P}(R) \to \{0,1\}$ of $\mathcal{LC}^R$ assigning 0 or 1 to each bit of the code, determines a family of subsets of R, $\mathcal{F}_E = E^{-1}(1) \subseteq \mathbf{P}(R)$. We do not distinguish between the codeword and the family determined by it. Thus we may say that the codeword encoding an element $e \in R$, is

$$\mathcal{F}_e = \{F \in \mathbf{P}(R) \ \mid \ F \ni e\} \ .$$

Background. The long-code was introduced in [BGS98], and utilized for obtaining numerous hardness results (some of which are tight) for approximation problems, such as Max-Cut, Max-2-Sat, Vertex-Cover, Max-3-Sat, Max-Lin-Eq, Max-Clique, Chromatic-Number. The scheme of [BGS98, Hås99, Hås97], is roughly as follows. Given a gap-SAT instance Ψ, first apply the parallel repetition lemma of [Raz98]. This, for any parameter k, results in a new gap-SAT problem, which we refer to as $\mathsf{Par}[\Psi, k]$, over (non-binary) variables of two types X and Y, and in which each local-constraint depends on one variable of X and one of Y. $\mathsf{Par}[\Psi, k]$ has a satisfying assignment if Ψ has one, and, otherwise, not even an arbitrarily small (exponentially small in k) fraction of the constrains can be satisfied.

In the next step, one applies a version of the Composition technique of [AS92], as proposed in [BGS98], to this specific setting, namely replaces each of the variables of X and Y, with a set of variables representing its long-code, and each local-constraint $\varphi(x, y) \in \mathsf{Par}[\Psi, k]$ with constraints that verify (i) *inner-consistency*, namely that the encoding of each variable is in some sense not far from a legal codeword (in fact, from a small list of acceptable codewords), and (ii) *outer-consistency*, namely that x and y's encoded values satisfy $\varphi(x, y)$.

Consider, for example, Håstad's linearity test. Given an encoding $\mathcal{F} \subset \mathbf{P}(R)$, consider the following random process. Choose two random subsets $F_1, F_2 \in \mathbf{P}(R)$, and a third subset $H \in \mathbf{P}(R)$ by taking each $e \in R$ to be in H independently with probability ϵ. Now, accept only if an even number (0 or 2) of the three subsets $F_1, F_2, F_1 \Delta F_2 \Delta H$ are in $\mathcal{F}$. If $\mathcal{F}$ is the true long-code of an element $e \in R$, this test accepts with probability $1 - \epsilon$. Moreover, one can prove, using extensive Fourier analysis, that if this test accepts with probability $\frac{1}{2} + \epsilon$, then $\mathcal{F}$ must be close to a true long-code.

The distribution according to which the subsets F_1, F_2 were chosen is uniform, implicitly implying that their size is roughly $\frac{1}{2} \cdot |R|$, except for a negligible fraction. The third subset, H, was chosen according to a distribution that highlights ϵ-sized subsets. We generalize these distributions, as follows.

The p-Biased Long-Code. Let us consider distributions that highlight subsets of size roughly $p \cdot |R|$. One such natural class of distributions, that highlight subsets whose size is roughly $p \cdot |R|$, is the *p–product-distribution* over $\mathbf{P}(R)$, denoted μ_p^R, where, independently for each element $e \in R$, e is in a set with probability p and out of it with probability $1 - p$. More precisely,

Definition 2.2 (Product-Distribution) *Let $0 < p < 1$. μ_p^R is a distribution over $\mathbf{P}(R)$ according to which, every subset $F \in \mathbf{P}(R)$ occurs with the following probability:*

$$\mu_p^R(F) \stackrel{def}{=} p^{|F|} \cdot (1-p)^{|R|-|F|}$$

In some cases, when the set R is clear from the context, we may omit R and refer to $\mu_p^R(F)$ simply as $\mu_p(F)$.

For $p = \frac{1}{2}$, μ_p is simply the uniform distribution. For other values of p, this distribution highlights sets whose cardinality is roughly $p \cdot |R|$, and turns out to be useful especially for $p < \frac{1}{2}$. Let us now introduce the p-biased long-code,

Definition 2.3 (The p-Biased Long-Code) *The* p-biased long-code *over R, denoted $\mathcal{LC}_p^R = \left\langle \mathbf{P}(R), \mu_p^R \right\rangle$, assigns the distribution μ_p to $\mathcal{LC}^R$.*

Motivation and Overview. Let us give some brief motivation for the analysis that follows. Aiming at proving hardness of approximating Minimum Vertex-Cover, we will construct, in future chapters, a weighted graph, whose vertices are partitioned into blocks, the vertices in each block corresponding to subsets $F \in \mathbf{P}(R)$ of the long-code of R. An independent set in this graph would correspond, in each block, to a family $\mathcal{F} \subseteq \mathbf{P}(R)$ obeying some combinatorial properties, supposedly encoding an element $e \in R$. In Section 2.1 we proceed to deduce, relying on some theorems from the field of influences of variables on Boolean functions, that such a family $\mathcal{F}$ distinguishes a *core*, namely, a small set $C \subset R$ of elements of R that are, in a sense, permissible decodings of it. In Section 2.2, we will show, that if $\mathcal{F}$ is also of large weight according to μ_p, and if it is *intersecting*, it must then distinguish, in a specific sense to be defined, one element in its core. This element will be important for asserting outer-consistency, as it will consequently be shown to be consistent with the distinguished elements of other encodings.

2.1 A Family's Core

Let $\mathcal{F} \subset \mathbf{P}(R)$ be a family of subsets of R. We would be interested in finding when this family is, in a sense, close to an encoding of an element $e \in R$. In fact, we would be satisfied in finding a small set of permissible elements in R, henceforth referred to as a core, such that $\mathcal{F}$ is roughly a combination of the codewords of these values.

A family of subsets $\mathcal{F} \subset \mathbf{P}(R)$ is said to be *determined* by $C \subset R$, if a subset $F \in \mathbf{P}(R)$ is determined to be in or out of $\mathcal{F}$ only according to its intersection with C (no matter whether other elements are in or out of F). Formally, $\mathcal{F}$ is determined by C if,

$$\{F \mid F \cap C \in \mathcal{F}\} = \mathcal{F}$$

Denote by $\mathcal{F}_1 \sqcup \mathcal{F}_2$ the family consisting of the pairwise union of all subsets of $\mathcal{F}_1$ with all those of $\mathcal{F}_2$,

$$\mathcal{F}_1 \sqcup \mathcal{F}_2 \overset{def}{=} \{F_1 \cup F_2 \mid F_1 \in \mathcal{F}_1, F_2 \in \mathcal{F}_2\} \ .$$

If $C \subset R$ determines $\mathcal{F}$, then there is a family $\mathcal{F}_C \subseteq \mathbf{P}(C)$, such that $\mathcal{F} = \mathcal{F}_C \sqcup \mathbf{P}(R \setminus C)$.

A given family $\mathcal{F}$, may not be determined by any small set C. However, there might be another family $\mathcal{F}'$, that is determined by some small set C, and that approximates $\mathcal{F}$ quite accurately, up to some δ:

Definition 2.4 (Core) *A set $C \subseteq R$ is said to be a (δ, p)-core of the family $\mathcal{F} \subseteq \mathbf{P}(R)$, if there exists $\mathcal{F}_C \subseteq \mathbf{P}(C)$ such that $\mu_p(\mathcal{F} \bigtriangleup (\mathcal{F}_C \sqcup \mathbf{P}(R \setminus C))) < \delta$ (where $\bigtriangleup$ denotes the symmetric difference between two families).*

As to the family of subsets that best approximates $\mathcal{F}$ on its core, it consists of the subsets $F \in \mathbf{P}(C)$ whose extension to R intersects more than half of $\mathcal{F}$,

$$\left\{ F \in \mathbf{P}(C) \;\middle|\; \Pr_{F' \in \mu_p^{R \setminus C}} \left[F \cup F' \in \mathcal{F} \right] > \frac{1}{2} \right\} .$$

Consider the *Core-Family*, defined as the family of all subsets $F \in \mathbf{P}(C)$, for which $\frac{3}{4}$ of their extension to R, i.e. $\frac{3}{4}$ of $\{F\} \sqcup \mathbf{P}(R \setminus C)$, reside in $\mathcal{F}$:

Definition 2.5 (Core-Family) *For a set of elements $C \subset R$, define,*

$$\langle \mathcal{F} \rangle_C \stackrel{def}{=} \left\{ F \in \mathbf{P}(C) \;\middle|\; \Pr_{F' \in \mu_p^R} \left[F' \in \mathcal{F} \mid F' \cap C = F \right] > \frac{3}{4} \right\}$$

By simple averaging, it turns out that if C is a (δ, p)-core for $\mathcal{F}$, this family approximates $\mathcal{F}$ almost as well as the best family C.

Lemma 2.1 *If C is a (δ, p)-core of $\mathcal{F}$, then $\mu_p^C\left(\langle \mathcal{F} \rangle_C \right) \geq \mu_p^R(\mathcal{F}) - 3\delta$.*

Proof: For each subset $F \in \mathbf{P}(C)$, consider

$$x_F \stackrel{def}{=} \Pr_{F' \in \mu_p} \left[F' \in \mathcal{F} \mid F' \cap C = F \right] .$$

A subset F is in $\langle \mathcal{F} \rangle_C$ iff $x_F > \frac{3}{4}$. Let $\mathcal{F}_C \subset \mathbf{P}(C)$ be the family that best approximates $\mathcal{F}$ on C, namely the family of subsets $F \in \mathbf{P}(C)$ for which $x_F > \frac{1}{2}$. By the definition of a (δ, p)-core, it follows

$$\mu_p^R(\mathcal{F}_C \sqcup \mathbf{P}(R \setminus C) \;\triangle\; \mathcal{F}) < \delta \,,$$

and we will prove that

$$\mu_p^R(\langle \mathcal{F} \rangle_C \sqcup \mathbf{P}(R \setminus C) \;\triangle\; \mathcal{F}) < 3\delta \,.$$

Replacing $\mathcal{F}_C$ by $\langle \mathcal{F} \rangle_C$ the symmetric difference from $\mathcal{F}$ is affected on those subsets F for which $\frac{1}{2} < x_F \leq \frac{3}{4}$, i.e. subsets $F \in \mathcal{F}_C \setminus \langle \mathcal{F} \rangle_C$. Observe that the relative contribution of each such subset to the symmetric difference from $\mathcal{F}$ increases from $(1 - x_F)$ to x_F. Since $\frac{1}{2} < x_F \leq \frac{3}{4}$, the symmetric difference is tripled at most. ■

Influence and Sensitivity

Understanding the conditions for family of subsets to have a small core, has been pursued, from a different perspective, for some years. This has to do with the probability of every element $e \in R$ to take subsets in or out of $\mathcal{F}$ when flipped, which is referred to as the *influence* of that element. This notion, and its relations with various properties of $\mathcal{F}$, have been the subject of an extensive analysis [BOL89, KKL88, Fri98]. Let us now introduce this notion and assert some theorems to be available for good use later.

Assume a family of subsets $\mathcal{F} \subseteq \mathbf{P}(R)$. The *influence* of an element $e \in R$,

$$\textbf{influence}_p^e(\mathcal{F}) \overset{def}{=} \Pr_{F \in \mu_p} [\text{ exactly one of } F \cup \{e\},\ F \setminus \{e\} \text{ is in } \mathcal{F}]$$

The *average sensitivity* of $\mathcal{F}$ with respect to μ_p, denoted $\mathbf{as}_p(\mathcal{F})$, is the sum of the influences of all elements in R,

$$\mathbf{as}_p(\mathcal{F}) \overset{def}{=} \sum_{e \in R} \textbf{influence}_p^e(\mathcal{F})$$

The name average-sensitivity is derived from the following. The *sensitivity* of a subset $F \in \mathcal{F}$ is the number of elements whose removal from or addition to F takes F in or out of $\mathcal{F}$:

$$|\{e \in R \mid \text{ exactly one of } F \cup \{e\},\ F \setminus \{e\} \text{ is in } \mathcal{F}\}|$$

The average sensitivity of $\mathcal{F}$ with respect to μ_p is equal to (hence the name) the average, according to μ_p, of the sensitivity of all subsets in $\mathbf{P}(R)$:

Proposition 2.2 *Let $\mathcal{F} \subseteq \mathbf{P}(R)$.*

$$\mathbf{as}_p(\mathcal{F}) = |R| \cdot \Pr_{F \in \mu_p, e \in_R R} [\textit{ exactly one of } F \cup \{e\},\ F \setminus \{e\} \textit{ is in } \mathcal{F}]$$

Proof: These are just two different ways to sum up the same set of events. ■

A truly fundamental relation between the average sensitivity of a family $\mathcal{F} \subseteq \mathbf{P}(R)$ and the size of its (δ, p)-core is the following theorem of Friedgut [Fri98]:

Theorem 2.3 (Friedgut) *Let $0 < p < 1$ be some bias, and $\delta > 0$ be any approximation parameter. Consider any family $\mathcal{F} \subset \mathbf{P}(R)$, and let $k = \mathbf{as}_p(\mathcal{F})$. There exists a function $\mathsf{h}(p, \delta, k) \leq (c_p)^{k/\delta}$, where c_p is a constant[1] depending only on p, such that $\mathcal{F}$ has a (δ, p)-core C, with $|C| \leq \mathsf{h}(p, \delta, k)$.* ■

Hence, the number of elements that are necessary in order to approximate $\mathcal{F}$ up to δ depends only on δ and the average sensitivity of $\mathcal{F}$. In particular, if a family $\mathcal{F}$ has low (say, constant) average sensitivity, then it has a (δ, p)-core whose size is merely exponential in $\frac{1}{\delta}$, and is independent of $|R|$.

The next step would be to find sufficient conditions for a family to have low average sensitivity. As it turns out, this is the case with *monotone families* (defined below), assuming we allow some slight shifting of p.

Definition 2.6 (Monotone Family) *A family of subsets $\mathcal{F} \subseteq \mathbf{P}(R)$ is* monotone *if for every $F \in \mathcal{F}$, for all $F' \supset F$, $F' \in \mathcal{F}$.*

[1] It follows directly from Friedgut's proof that c_p can be taken as a continuous function of p.

Such a family is sometimes called in the literature an 'upset'.

The *monotone closure* $\bar{\mathcal{F}}$ of a family $\mathcal{F}$ is defined to be the family of all subsets containing a subset from $\mathcal{F}$, $\bar{\mathcal{F}} \overset{def}{=} \{F \cup F' \mid F \in \mathcal{F}\}$.

Being monotone restricts a family in certain ways, forcing it, for example, to have relatively more large subsets than it does small subsets. This can be formalized as follows,

Proposition 2.4 *For a monotone family $\mathcal{F} \subseteq \mathbf{P}(R)$, $\mu_p(\mathcal{F})$ is a monotone non-decreasing function of p.*

Proof: Assume $R = [n]$. For a subset $F \in \mathbf{P}([n])$ denote

$$F_{\leq i} \overset{def}{=} F \cap [1, i] \quad \text{and} \quad F_{>i} \overset{def}{=} F \cap [i+1, n]$$

and consider, for $0 \leq i \leq n$, the hybrid distribution, where the first i elements are chosen with bias p and the others are chosen with bias $q > p$,

$$\mu_{p,i,q}(F) \overset{def}{=} p^{\left|F_{\leq i}\right|} \cdot (1-p)^{i-\left|F_{\leq i}\right|} \cdot q^{|F_{>i}|} \cdot (1-q)^{n-i-|F_{>i}|} .$$

Observe that

$$\forall\, 0 \leq i \leq n \quad \mu_{p,i,q}(\mathcal{F}) \geq \mu_{p,i+1,q}(\mathcal{F})$$

therefore $\mu_q(\mathcal{F}) = \mu_{p,0,q}(\mathcal{F}) \geq \mu_{p,n,q} = \mu_p(\mathcal{F})$. ■

Interestingly, for monotone families, the rate at which μ_p increases with p, is exactly equal to the average-sensitivity:

Theorem 2.5 (Russo-Margulis Identity [Mar74, Rus82]) *Let $\mathcal{F} \subseteq \mathbf{P}(R)$ be a monotone family. Then,*

$$\frac{d\mu_p(\mathcal{F})}{dp} = \mathbf{as}_p(\mathcal{F})$$

Proof: For a subset $F \in \mathbf{P}(R)$ write

$$\mu_p(F) = \prod_{e \in R} \mu_p^e(F), \qquad \text{for} \quad \mu_p^e(F) = \begin{cases} p & e \in F \\ 1-p & e \notin F \end{cases} \tag{$*$}$$

Observe that

$$\mathbf{influence}_p^e(\mathcal{F}) = \sum_{F \in \mathcal{F}} \left(\frac{d\mu_p^e(F)}{dp} \cdot \prod_{e' \neq e} \mu_p^{e'}(F) \right)$$

Differentiating $(*)$ according to p, and summing over all $F \in \mathcal{F}$, we get

$$\frac{d\mu_p(\mathcal{F})}{dp} = \sum_{e \in R} \mathbf{influence}_p^e(\mathcal{F}) = \mathbf{as}_p(\mathcal{F})$$

■

The average sensitivity of a monotone family, depends only on the number of subsets of every given size,

Proposition 2.6 *Let $\mathcal{F} \subseteq \mathbf{P}(R)$ be monotone, let $\mathcal{F}_k = \{F \in \mathcal{F} \mid |F| = k\}$, $n = |R|$,*

$$\mathbf{as}_p(\mathcal{F}) = \sum_k p^k(1-p)^{n-k}\,|\mathcal{F}_k| \cdot \left(\frac{1}{p} \cdot k - \frac{1}{1-p} \cdot (n-k)\right)$$

Proof: For a pair of subsets $F, F \setminus \{i\} \in \mathbf{P}(R)$, if exactly one is in $\mathcal{F}$ then by the monotonicity of $\mathcal{F}$, $F \in \mathcal{F}$ and $F \setminus \{i\} \notin \mathcal{F}$. By Proposition 2.2,

$$\mathbf{as}_p(\mathcal{F}) = \sum_{i \in F \in \mathcal{F}, F \setminus \{i\} \notin \mathcal{F}} \mu_p(F) + \mu_p(F \setminus \{i\}) \ .$$

By writing $\mu_p(F) + \mu_p(F \setminus \{i\}) = \frac{1}{p} \cdot \mu_p(F)$ and regrouping, this becomes,

$$\begin{aligned} &= \sum_{F \in \mathcal{F}} \left(\sum_{i \in F,\ F \setminus \{i\} \notin \mathcal{F}} \frac{1}{p} \cdot \mu_p(F) \right) = \frac{1}{p} \cdot \sum_{F \in \mathcal{F}} \left(\mu_p(F)\,|F| - \sum_{i \in F,\ F \setminus \{i\} \in \mathcal{F}} \mu_p(F) \right) \\ &= \frac{1}{p} \cdot \sum_{F \in \mathcal{F}} \mu_p(F)\,|F| - \frac{1}{p} \cdot \sum_{i \in F \in \mathcal{F},\ F \setminus \{i\} \in \mathcal{F}} \mu_p(F) \end{aligned}$$

Note now that both sums are taken over subsets inside $\mathcal{F}$. Rename in the second sum $F_1 = F \setminus \{i\}$, and note that each subset F_1 appears in the second sum exactly $|R \setminus F_1|$ times. By writing $\mu_p(F) = \mu_p(F_1) \cdot \frac{p}{1-p}$, the second sum equals $\sum_{F_1 \in \mathcal{F}} \mu_p(F_1) \cdot \frac{|R \setminus F_1|}{1-p}$ and together

$$= \sum_{F \in \mathcal{F}} \mu_p(F) \cdot \left(\frac{|F|}{p} - \frac{|R \setminus F|}{1-p} \right) \ .$$

Summing subsets in $\mathcal{F}$ according to their size k, gives the claim. ■

This identity shows that 'slices' of $\mathcal{F}$, $\mathcal{F}_k$, with $k > p \cdot n$, yield a positive contribution to the average sensitivity, while slices with $k < p \cdot n$ contribute negatively. Hence, the *threshold family*, whose positive slices are completely full, and whose negative slices are completely empty,

$$\mathcal{F}_{\geq p} \overset{def}{=} \{F \in \mathbf{P}(R) \mid |F| \geq p \cdot |R|\}$$

has maximal average sensitivity $\mathbf{as}_p(\mathcal{F}) = \Theta(\sqrt{|R|})$. One may note that this is still well below the maximal average sensitivity of any general family (which is $|R|$, attained by the parity family) .

An interesting thing to notice for the threshold family $\mathcal{F}_{\geq p}$, is that although its average sensitivity is high, considering this family with a slightly shifted p, say $p + \epsilon$ or $p - \epsilon$ for some small but constant $\epsilon > 0$, makes this family be approximately full $\mu_{p+\epsilon}(\mathcal{F}_{\geq p}) \approx 1$ or approximately empty $\mu_{p-\epsilon}(\mathcal{F}_{\geq p}) \approx 0$, and in both cases the average sensitivity according to both $\mu_{p+\epsilon}$ and $\mu_{p-\epsilon}$ is almost zero.

In fact, a similar argument holds for *every* monotone family. As we gradually increase p, the average sensitivity $\mathbf{as}_p(\mathcal{F})$ of a monotone family $\mathcal{F}$, although possibly remaining non-zero, cannot be very high for too long:

Proposition 2.7 *Let $\mathcal{F} \subseteq \mathbf{P}(R)$ be a monotone family, and let $0 \leq p < p + \epsilon \leq 1$. There must be some $q \in (p, p+\epsilon)$ such that*

$$\mathbf{as}_q(\mathcal{F}) \leq \frac{1}{\epsilon}$$

Proof: With the above identity, and a standard application of Lagrange's Mean-Value Theorem, there exists some $q \in (p, p+\epsilon)$,

$$\mathbf{as}_q(\mathcal{F}) = \frac{d\mu_q(\mathcal{F})}{dq} = \frac{\mu_{p+\epsilon}(\mathcal{F}) - \mu_p(\mathcal{F})}{\epsilon} \leq \frac{1}{\epsilon}$$

■

We have now reached the main point of this discussion. A monotone family $\mathcal{F}$, supposedly representing an encoding with the p-biased long code of an element in R, always has low average sensitivity for some value of $q \in (p, p+\epsilon)$. For this q we can apply Friedgut's Lemma to deduce a small core $C \subset R$, $|C| = O(1)$, for $\mathcal{F}$, on which it is well-approximated according to μ_q. The elements in this core would serve as a set of permissible values, that are the 'decoding' of $\mathcal{F}$, in the rest of the proof. That these decoded values indeed represent $\mathcal{F}$, and that consistency of families $\mathcal{F}_1$ and $\mathcal{F}_2$ constitute some form of consistency of their cores C_1 and C_2, is the task we face in the next chapter.

Let us conclude this section with an easy proposition, to be used later on, showing that, if $T \subset R$ is a set of elements of tiny influence in a monotone family $\mathcal{F} \subset \mathbf{P}(R)$, one has to remove only a small fraction of $\mathcal{F}$ to make it completely independent of T:

Proposition 2.8 *Let $\mathcal{F} \subset \mathbf{P}(R)$ be monotone, and let $T \subset R$ be such that for all $e \in T$, $\mathbf{influence}^e_p(\mathcal{F}) < \eta$. Let*

$$\mathcal{F}' = \{F \in \mathcal{F} \mid F \setminus T \in \mathcal{F}\}$$

then,

$$\mu^R_p(\mathcal{F} \setminus \mathcal{F}') < |T| \cdot \eta \cdot p^{-|T|}$$

Proof: Let

$$\mathcal{F}'' = \{F \in \mathbf{P}(R \setminus T) \mid F \cup T \in \mathcal{F} \text{ but } F \notin \mathcal{F}\} .$$

A set $F \in \mathcal{F}''$ contributes at least $\mu^{R\setminus T}_p(F) \cdot p^{|T|}$ to the influence of at least one element $e \in T$, so $\mu^{R\setminus T}_p(\mathcal{F}'') < |T| \cdot \eta \cdot p^{-|T|}$. The proof is complete noting that,

$$\mathcal{F} \setminus \mathcal{F}' \subseteq \mathcal{F}'' \sqcup \mathbf{P}(T)$$

■

2.2 Maximal Intersecting Families

We have seen in the previous section, that a monotone family distinguishes a small core of elements, that almost determine it completely. In this section we will show that a monotone family that is of large enough weight, and is also *intersecting*, must exhibit one *distinguished* element in its core. This element is a stricter 'decoding' of the family than is the core, and will consequently serve for establishing outer-consistency, i.e. consistency between families encoding distinct variables.

Consider the encoding of an element $e \in R$, according to the p-biased long-code, namely the family of all subsets containing e, $\mathcal{F}_e = \{F \in \mathbf{P}(R) \mid F \ni e\}$. This family is both monotone, as defined above, and intersecting, defined next,

Definition 2.7 (t-Intersecting Family) *A family $\mathcal{F} \subset \mathbf{P}(R)$ is said to be* t-intersecting*, for $t \geq 1$, if*

$$\forall F_1, F_2 \in \mathcal{F}, \quad |F_1 \cap F_2| \geq t\,.$$

If $t = 1$ such a family is referred to simply as intersecting*.*

The following is a natural generalization for a pair of families,

Definition 2.8 (Cross-Intersecting) *Two families $\mathcal{F}_1, \mathcal{F}_2 \subseteq \mathbf{P}(R)$ are* cross-intersecting *if for every $F_1 \in \mathcal{F}_1$ and $F_2 \in \mathcal{F}_2$, $F_1 \cap F_2 \neq \phi$.*

Two families cannot be too large and still remain cross-intersecting,

Proposition 2.9 *For any bias parameter $p \leq \frac{1}{2}$, two families of subsets $\mathcal{F}_1, \mathcal{F}_2 \subseteq \mathbf{P}(R)$, for which $\mu_p(\mathcal{F}_1) + \mu_p(\mathcal{F}_2) > 1$ are not cross-intersecting.*

Proof: We can assume that $\mathcal{F}_1, \mathcal{F}_2$ are monotone, as their monotone closures must also be cross-intersecting. Since μ_p, for a monotone family, is non-decreasing with respect to p, it is enough to prove the claim for $p = \frac{1}{2}$. If for all $F \in \mathbf{P}(R)$ contained in both families – that is, so that $F \in \mathcal{F}_1$ and $F \in \mathcal{F}_2$ – it were the case that its complement $F^{\mathsf{c}} = R \setminus F$ would be contained in none of the families – namely, $F^{\mathsf{c}} \notin \mathcal{F}_1, F^{\mathsf{c}} \notin \mathcal{F}_2$ – the sum of sizes would be at most 1. There must therefore be one such pair, F and F^{c}, contained one in $\mathcal{F}_1$ and the other in $\mathcal{F}_2$. ■

It is now easy to prove that if $\mathcal{F}$ is monotone and intersecting, then the same holds for the core-family $\langle \mathcal{F} \rangle_C$ that is (see Definition 2.5) the threshold approximation of $\mathcal{F}$ on its core C,

Proposition 2.10 *Let $\mathcal{F} \subseteq \mathbf{P}(R)$, and let C be a (δ, p)-core of $\mathcal{F}$.*

- *If $\mathcal{F}$ is monotone then $\langle \mathcal{F} \rangle_C$ is monotone.*
- *If $\mathcal{F}$ is intersecting, and $p \leq \frac{1}{2}$, then $\langle \mathcal{F} \rangle_C$ is intersecting.*

Proof: The first assertion is immediate. For the second assertion, assume a pair of non-intersecting subsets $F_1, F_2 \in \langle \mathcal{F} \rangle_C$ and observe that the families

$$\{F \in \mathbf{P}(R \setminus C) \mid F \cup F_1 \in \mathcal{F}_1\} \qquad \text{and} \qquad \{F \in \mathbf{P}(R \setminus C) \mid F \cup F_2 \in \mathcal{F}_2\}$$

both have weight $> \frac{3}{4}$, and by Proposition 2.9, cannot be cross intersecting. ■

It can be proven that cross-intersecting families must have intersecting cores, unless the families are of negligible size. This alone does not provide sufficient consistency for our construction, as one family can appear to be consistent with other families, by exhibiting a different element in its core for consistency with different families.

Note however that any family cross-intersecting $\mathcal{F}_e$, must have $e \in R$ in its core. Simplistically, we can imitate this property showing that a non-2-intersecting family $\mathcal{F}$, must have a distinguished element (an element that is the exact intersection of a pair of subsets in $\mathcal{F}$), and that this element must reside in the cores of families consistent with $\mathcal{F}$.

Definition 2.9 (Distinguished Element) *For a monotone and intersecting family $\mathcal{F} \subseteq \mathbf{P}(R)$, an element $e \in R$ is said to be* distinguished *if there exist $F^\flat, F^\sharp \in \mathcal{F}$ such that*

$$F^\flat \cap F^\sharp = \{e\}$$

Clearly, an intersecting family has a distinguished element if and only if it is not 2-intersecting. We next establish a weight criterion for an intersecting family to have a distinguished element, summarized as follows:

For each $p < 0.4$, we define $p^\bullet$ to be,

Definition 2.10

$$\forall p < 0.4, \qquad p^\bullet \overset{def}{=} \max(p^2,\, 4p^3 - 3p^4)$$

This maps each p to the size of the maximal 2-intersecting family, according to μ_p, as asserted by the following lemma,

Lemma 2.13 *If $\mathcal{F} \subset \mathbf{P}(R)$ is monotone and 2-intersecting, then $\mu_p(\mathcal{F}) \leq p^\bullet$, provided $p < 0.4$.*

For a proof of the above we venture into the field of extremal set theory, where maximal intersecting families have been studied for some time. This beautiful study began with a paper of Erdős, Ko, and Rado [EKR61], that has seen various extensions and generalizations. The lemma above is a generalization to μ_p of what is known as the Complete Intersection Theorem for finite sets, that was proven by [AK97].

This section is broken into two subsections. In the first subsection we present the theorem of Ahlswede and Khachatrian [AK97] (Theorem 2.11), and prove Lemma 2.13. In the second subsection, we review some known results pertaining to maximal intersecting families, generalizing some of them for μ_p, leading up to a proof of the upper bound of p for the weight of a monotone intersecting family.

2.2.1 2-Intersecting Families

For the following, assume $R = [n]$. In the field of extremal set theory, the focus is on families of subsets of a given size k. Denote the k-th slice of $\mathbf{P}([n])$ by

$$\binom{[n]}{k} \overset{def}{=} \{F \in \mathbf{P}([n]) \mid |F| = k\}.$$

The Erdős-Ko-Rado Theorem [EKR61], states that given $k, t \geq 1$, then for large enough n, the maximal t-intersecting family $\mathcal{F} \subseteq \binom{[n]}{k}$ is obtained by taking all subsets that contain t fixed elements, thus $|\mathcal{F}| \leq \binom{n-t}{k-t}$.

As the measure μ_p is concentrated on subsets of size near $k = p \cdot n$, bounds for t-intersecting families $\mathcal{F} \subset \binom{[n]}{k}$, would translate to bounds for the weight – $\mu_p(\mathcal{F})$ – of monotone t-intersecting families. The Erdős-Ko-Rado Theorem does not suffice, because as n increases, we need to maintain a constant proportion between n and k. Frankl [Fra78], investigated the full range of values for n and k, and conjectured that the maximal t-intersecting family is always one of

$$\mathcal{A}_{i,t} \overset{def}{=} \{F \in \mathbf{P}([n]) \mid F \cap [1, t+2i] \geq t+i\}$$

Partial versions of this conjecture have been proven by [Fra78, FF91, Wil84], and the complete intersection theorem for finite sets was finally proven by Ahlswede and Khachatrian [AK97],

Theorem 2.11 ([AK97]) *Let $\mathcal{F} \subseteq \binom{[n]}{k}$ be t-intersecting. Then,*

$$|\mathcal{F}| \leq \max_{0 \leq i \leq \frac{n-t}{2}} \left|\mathcal{A}_{i,t} \cap \tbinom{[n]}{k}\right|$$

Our analysis requires the extension of this statement to families of subsets that are not restricted to one size k. We restrict our attention to $t = 2$, and denote $\mathcal{A}_i \overset{def}{=} \mathcal{A}_{i,2}$.

Lemma 2.12 *Let $\mathcal{F} \subset \mathbf{P}(R)$ be monotone and 2-intersecting. For any $p < \frac{1}{2}$,*

$$\mu_p(\mathcal{F}) \leq \max_i \{\mu_p(\mathcal{A}_i)\}.$$

As a direct corollary of this lemma, we obtain the lemma that was stated in the beginning of this section.

Lemma 2.13 *If $\mathcal{F} \subset \mathbf{P}(R)$ is monotone and 2-intersecting, then $\mu_p(\mathcal{F}) \leq p^\bullet$, provided $p < 0.4$.*

Let us first prove the lemma,

Proof: Assume $\mathcal{F}_0 \subset \mathbf{P}([n_0])$ contradicts the claim, denote $\mu = \max_i(\mu_p(\mathcal{A}_i))$, and let $a = \mu_p(\mathcal{F}_0) - \mu > 0$. Now consider $\mathcal{F} = \mathcal{F}_0 \sqcup \mathbf{P}([n] \setminus [n_0])$ for $n > n_0$ large enough, to be determined below. Clearly, for any $n \geq n_0$, $\mu_p^{[n]}(\mathcal{F}) = \mu_p^{[n_0]}(\mathcal{F}_0)$, and $\mathcal{F}$ is 2-intersecting. Consider the 'slices' near $p \cdot n$, (let $\theta < \frac{1}{2} - p$)

$$S \overset{def}{=} \{k \in \mathbb{N} \mid |k - p \cdot n| \leq \theta \cdot n\}$$

and for every $k \in S$, denote by $\mathcal{F}_k = \mathcal{F} \cap \binom{[n]}{k}$. We will show that since most of $\mathcal{F}$'s weight is derived from $\cup_{k \in S}\mathcal{F}_k$, there must be at least one $\mathcal{F}_k$ that contradicts Theorem 2.11. Indeed,

$$\mu + a = \mu_p(\mathcal{F}) = \sum_{k \in S} p^k(1-p)^{n-k} \cdot |\mathcal{F}_k| + o(1)$$

Hence there exists $k \in S$ for which $\frac{|\mathcal{F}_k|}{\binom{[n]}{k}} \geq \mu + \frac{1}{2}a$. We have left to see that $\mu \cdot \binom{n}{k}$ is close enough to $\max_i(|\mathcal{A}_i \cap \binom{[n]}{k}|)$. This follows from usual tail bounds, and is sketched as follows. Subsets in $\binom{[n]}{k}$ for large enough i (depending only on $\frac{k}{n}$ but not on k or n), have roughly $\frac{k}{n} \cdot (2i+2)$ elements in the set $[1, 2i+2]$. Moreover, the subsets in $\mathcal{A}_i$ have at least $i+2$ elements in $[1, 2i+2]$, thus are very few (compared to $\binom{n}{k}$), because $\frac{i+2}{2i+2} > \frac{1}{2} > p + \theta \geq \frac{k}{n}$. In other words, there exists some constant $C_{p+\theta,\mu}$, for which $\left|\mathcal{A}_i \cap \binom{[n]}{k}\right| < \mu \cdot \binom{n}{k}$ for all $i \geq C_{p,\mu}$ as long as $\frac{k}{n} \leq p + \theta$.

Additionally, for every $i < C_{p,\mu}$, taking n to be large enough we have

$$\forall k \in S, \qquad \frac{\left|\mathcal{A}_i \cap \binom{[n]}{k}\right|}{\binom{n}{k}} = \mu_{\frac{k}{n}}(\mathcal{A}_i) + o(1) = \mu_p(\mathcal{A}_i) + o(1) < \mu + o(1)$$

where the first equality follows from a straightforward computation. ■

We next proceed to derive Lemma 2.13 from the above,

Proof: Define a sequence $p_0 < p_1 < \ldots$, where $p_i \overset{def}{=} \frac{i}{2i+1}$. We will show that these are the points where the maximum switches from $\mathcal{A}_i$ to $\mathcal{A}_{i+1}$. More accurately, we will show for all $i \geq 0$,

$$\forall p \in (p_i, p_{i+1}] \quad \max_j \{\mu_p(\mathcal{A}_j)\} = \mu_p(\mathcal{A}_i) \qquad (*)$$

This, together with Lemma 2.12, will complete our proof, as $p < p_{max} < 0.4 = p_2$ implies $\mu_p(\mathcal{F}) \leq \max(\mu_p(\mathcal{A}_0), \mu_p(\mathcal{A}_1)) = \max(p^2, 4p^3 - 3p^4) = p^\bullet$.

So we proceed to prove $(*)$. A subset $F \notin \mathcal{A}_i$ must intersect $[1, 2i+2]$ on at most $i+1$ elements. If additionally $F \in \mathcal{A}_{i+1}$ it must then contain $2i+3, 2i+4$. Thus,

$$\mu_p(\mathcal{A}_{i+1} \setminus \mathcal{A}_i) = \binom{2i+2}{i+1} \cdot p^{i+1}(1-p)^{i+1} \cdot p^2$$

Similarly,

$$\mu_p(\mathcal{A}_i \setminus \mathcal{A}_{i+1}) = \binom{2i+2}{i+2} \cdot p^{i+2}(1-p)^i \cdot (1-p)^2$$

We write,

$$\begin{aligned}\mu_p(\mathcal{A}_{i+1}) - \mu_p(\mathcal{A}_i) &= \mu_p(\mathcal{A}_{i+1} \setminus \mathcal{A}_i) - \mu_p(\mathcal{A}_i \setminus \mathcal{A}_{i+1}) \\ &= p^{i+2}(1-p)^{i+1}\binom{2i+2}{i+1}\left(p - (1-p)\frac{i+1}{i+2}\right)\end{aligned}$$

The sign of this difference is determined by $p - (1-p)\frac{i+1}{i+2}$. For a fixed $i \geq 0$, this expression goes from positive to negative passing through zero once at $p = \frac{i+1}{2i+3} = p_{i+1}$. Thus, the sequence $\{\mu_p(\mathcal{A}_j)\}_j$ is maximized at i for $p_i < p \leq p_{i+1}$. (It is increasing when $i \leq \frac{1-3p}{2p-1}$, and decreasing thereafter). ■

2.2.2 Intersecting Families

In this subsection we focus on bounding the maximal size of an intersecting family. This bound can be derived directly from Lemma 2.12 above, however, we present here another approach, that explains some of the structure of intersecting families. All the results in this subsection are either well known, or are easily derived from well-known results.

An intersecting family $\mathcal{F} \subset \mathbf{P}(R)$ never contains both F and its complement $F^c \overset{def}{=} R \setminus F$. For a monotone family this is a necessary as well as sufficient condition. Thus, for such $\mathcal{F}$, $\mu_{\frac{1}{2}}(\mathcal{F}) \leq \frac{1}{2}$. The same argument also implies $|\mathcal{F} \cap \binom{[n]}{k}| + |\mathcal{F} \cap \binom{[n]}{n-k}| \leq \binom{n}{k}$, an observation helpful for proving the following generalization for any $p \leq \frac{1}{2}$,

Lemma 2.14 *Let $\mathcal{F}$ be monotone and intersecting, and let $p \leq \frac{1}{2}$. Then*

$$\mu_p(\mathcal{F}) \leq p\,.$$

Thus, the legal encoding of $e \in R$, $\mathcal{F}_e$, is a monotone-intersecting family of maximum weight. Note that for $p > \frac{1}{2}$, the *majority* family (the family of all subsets whose size is more than $\frac{1}{2} \cdot R$) has weight almost 1, much higher than p.

For a proof of the above let us define some notions.

Shadows and the Kruskal-Katona Theorem. Assume $R = [n]$, and let $\mathcal{F} \subseteq \mathbf{P}([n])$. The (upper) shadow of $\mathcal{F}$, is defined to be

$$\partial\mathcal{F} \overset{def}{=} \{F \cup \{i\} \mid F \in \mathcal{F},\ i \notin F\}\ .$$

If $\mathcal{F} \subseteq \binom{[n]}{k}$, then $\partial\mathcal{F} \subseteq \binom{[n]}{k+1}$. Clearly, the class of families containing their own shadow, is exactly the class of monotone families; and the subsets in $\mathcal{F} \setminus \partial\mathcal{F}$ are the *minterms* of $\mathcal{F}$.

The l-th shadow of $\mathcal{F}$ is defined inductively to be $\partial^{(l)}\mathcal{F} \overset{def}{=} \partial(\partial^{(l-1)}\mathcal{F})$, and is simply the family of subsets that contain some subset of $\mathcal{F}$ and any l additional elements.

It is easy to see that a family $\mathcal{F} \subseteq \binom{[n]}{k}$ is intersecting, if and only if its $(n-2k)$-th shadow does not contain any set that is complementary to a set in $\mathcal{F}$. Bounding the size of maximal intersecting families gives rise to the following question: Given a family $\mathcal{F} \subseteq \binom{[n]}{k}$ of certain cardinality, what is the minimal size of its shadow? This question is answered by the Kruskal-Katona theorem, below.

An interesting and useful order when considering families of subsets (and in particular, intersecting or monotone families), is the *reverse lexicographic order* on $\binom{[n]}{k}$; defined by reversing the order that comes from interpreting the characteristic vector of each subset as a binary representation of an integer:

$$\forall F_1, F_2 \in \binom{[n]}{k}, \quad F_1 < F_2 \quad \text{if} \quad \sum_{i\in F_1} 2^{-i} > \sum_{i\in F_2} 2^{-i}$$

According to this order, we define

Definition 2.11 *Given $\mathcal{F} \subseteq \binom{[n]}{k}$, its* alignment $L(\mathcal{F}) \subseteq \binom{[n]}{k}$ *is the family consisting of the first $|\mathcal{F}|$ subsets in the lexicographic order. The alignment of $\mathcal{F} \subseteq \mathbf{P}([n])$, is defined by aligning each slice $\mathcal{F} \cap \binom{[n]}{k}$ separately, $L(\mathcal{F}) \overset{def}{=} \bigcup_k L(\mathcal{F} \cap \binom{[n]}{k})$*

It turns out that aligning a family 'compresses' it, and can only reduce the size of its shadow:

Theorem 2.15 ([Kru63, Kat68]) *For any $\mathcal{F} \subseteq \binom{[n]}{k}$, $\partial(L(\mathcal{F})) \subseteq L(\partial\mathcal{F})$.*

It follows in particular, since $|\partial\mathcal{F}| = |L(\partial\mathcal{F})|$, that $|\partial\mathcal{F}| \geq |\partial L(\mathcal{F})|$ thus $L(\mathcal{F})$ has minimal sized shadow. For a proof of this theorem, and an instructive exposition to this whole topic, see chapters 5,7,13 in [Bol86]. The following properties hold for $L(\mathcal{F})$,

Proposition 2.16 *Let $\mathcal{F} \subseteq \mathbf{P}(R)$.*

1. *If $\mathcal{F}$ is monotone, then $L(\mathcal{F})$ is monotone.*
2. *If $\mathcal{F}$ is monotone, then $\mathbf{as}_p(\mathcal{F}) = \mathbf{as}_p(L(\mathcal{F}))$.*
3. *$\mu_p(\mathcal{F}) = \mu_p(L(\mathcal{F}))$.*

Proof: $\mathcal{F}$ is monotone iff $\mathcal{F}_{k+1} \supseteq \partial(\mathcal{F}_k)$. Aligning these two families, and by theorem 2.15,

$$L(\mathcal{F}_{k+1}) \supseteq L(\partial\mathcal{F}_k) \supseteq \partial L(\mathcal{F}_k)$$

Both the average sensitivity $\mathbf{as}_p(\mathcal{F})$ in case $\mathcal{F}$ is monotone (see proposition 2.6), and the weight $\mu_p(\mathcal{F})$, depend only on the cardinality of each slice of $\mathcal{F}$, hence remain unchanged by aligning, because $\left|\mathcal{F} \cap \binom{[n]}{k}\right| = \left|L(\mathcal{F}) \cap \binom{[n]}{k}\right|$. ■

We must note that although for monotone families, $\mathbf{as}_p(\mathcal{F}) = \mathbf{as}_p(L(\mathcal{F}))$, and although the average sensitivity is the sum of the influences of the elements, it is certainly *not* the case that $L(\mathcal{F})$ leaves the influences of distinct elements unchanged. Finally, we have

Proposition 2.17 *If $\mathcal{F}$ is intersecting, then $L(\mathcal{F})$ is intersecting.*

Proof: If $\mathcal{F}$ is intersecting, then its monotone closure, $\bar{\mathcal{F}}$, is also intersecting, and of course $L(\bar{\mathcal{F}}) \supseteq L(\mathcal{F})$. Hence, it suffices to prove the proposition for $\mathcal{F}$ both monotone and intersecting.

In this case, $L(\mathcal{F})$ is also monotone by Proposition 2.16 above. For $L(\mathcal{F})$ to be intersecting, it is enough to show that for every F, not both F and $F^{\mathsf{c}} = [n] \setminus F$ are in $L(\mathcal{F})$. $\mathcal{F}$ is intersecting, so

$$|L(\mathcal{F}_k)| + |L(\mathcal{F}_{n-k})| = |\mathcal{F}_k| + |\mathcal{F}_{n-k}| \leq \binom{n}{k}$$

The proof is complete once we observe that the first m subsets in the lexicographic order on $\binom{[n]}{k}$, are the complements of the last m subsets in the lexicographic order on $\binom{[n]}{n-k}$, so the above inequality prevents two complement sets from being in $L(\mathcal{F})$. ■

Let us remark that if $\mathcal{F}$ is t-intersecting for $t > 1$, its alignment $L(\mathcal{F})$ is not necessarily t-intersecting, so these techniques do not carry over for bounding the size of general t-intersecting families. We now return to prove

Lemma 2.14 *Let $\mathcal{F}$ be a monotone and intersecting family, and let $p \leq \frac{1}{2}$. Then*

$$\mu_p(\mathcal{F}) \leq p\,.$$

Proof: By propositions 2.16 and 2.17, $L(\mathcal{F})$ is also monotone and intersecting, and $\mu_p(\mathcal{F}) = \mu_p(L(\mathcal{F}))$. Thus we assume w.l.o.g. that $\mathcal{F} = L(\mathcal{F})$. Define the following monotone and intersecting family,

$$\mathcal{F}^1 = \{F \in \mathbf{P}\left([n]\right) \mid F \ni 1\}$$

and note that $|\mathcal{F}^1_k| + \left|\mathcal{F}^1_{n-k}\right| = \binom{n-1}{k-1} + \binom{n-1}{n-k-1} = \binom{n}{k}$.

We rely on the easy fact that for any intersecting $\mathcal{F}$, $|\mathcal{F}_k| + |\mathcal{F}_{n-k}| \leq \binom{n}{k}$, as always only one of F and F^{c} can be in $\mathcal{F}$. Observe that if $\mu_p(\mathcal{F}) > p = \mu_p(\mathcal{F}^1)$, and since $p \leq \frac{1}{2}$, there must be some $k < n-k$ for which $|\mathcal{F}_k| > |\mathcal{F}^1_k|$. By $\mathcal{F} = L(\mathcal{F})$ we deduce that $\mathcal{F}_k \supsetneq \mathcal{F}^1_k$. By the monotonicity of $\mathcal{F}$, and by the Kruskal-Katona theorem (Theorem 2.15), we deduce that

$$\mathcal{F}_{n-k} \supseteq \partial^{(n-2k)}(\mathcal{F}_k) \supseteq \partial^{(n-2k)}(\mathcal{F}^1_k) = \mathcal{F}^1_{n-k}$$

so $|\mathcal{F}_k| + |\mathcal{F}_{n-k}| > |\mathcal{F}^1_k| + \left|\mathcal{F}^1_{n-k}\right| = \binom{n}{k}$, a contradiction. ■

Chapter 3

Reducing PCP to Vertex Cover

The aim of this chapter is to construct a graph, for which it is NP-hard to approximate the size of the smallest vertex cover. Our construction begins with a gap-SAT instance Φ, and transforms it into a graph $\mathcal{G}_{[c,s]}(\Phi)$ whose independent set is large $\alpha(\mathcal{G}_{[c,s]}(\Phi)) \geq c$, in case Φ is satisfiable (this is the *c*ompleteness of the reduction), and small $\alpha(\mathcal{G}_{[c,s]}(\Phi)) \leq s$, in case Φ is far from satisfiable (this is the *s*oundness of the reduction). The factor of hardness thus obtained for vertex cover, is $\frac{1-s}{1-c}$.

Weighted-Graphs. Our analysis is more naturally presented over weighted graphs, where the size of a set of vertices is the sum of their weights. Hardness results for these graphs easily translate to hardness for graphs with equal weight, see Appendix A.

A *weighted-graph* $G = (V, E, \Lambda)$ is an undirected graph with vertices V and edges E, and a probability distribution Λ over the vertices V. In other words, G is a graph with normalized weights. An *independent set* in G is a set $\mathcal{I} \subseteq V$ such that G restricted to $\mathcal{I}$ is the empty graph. Let us denote by $\alpha(G)$ the *maximum*, over all independent sets $\mathcal{I}$ in G, of $\Lambda(\mathcal{I})$. A *vertex-cover* of G is a set $S \subseteq V$ whose complement $V \setminus S$ is an independent set. Let us denote by $\overline{\alpha}(G)$ the *minimum*, over all vertex-covers S, of $\Lambda(S)$.

Our method for constructing $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ is very roughly – following the composition framework – to modify the gap-SAT instance Φ, and then apply the p-biased long-code over this modification. We will prove (in Chapter 4) that if Φ is satisfiable, then $\alpha(\mathcal{G}_{[p,p^\bullet]}(\Phi)) > p - \epsilon$, and if Φ is far from satisfiable, then $\alpha(\mathcal{G}_{[p,p^\bullet]}(\Phi)) < p^\bullet + \epsilon$.

Outline. In Section 3.1, we formally describe how to obtain our starting point – namely the gap-SAT problem – by applying the parallel repetition lemma of [Raz98] to the PCP theorem of [AS92, ALM$^+$92]. In Section 3.2 we describe the FGLSS-graph which will be used as a skeleton for our final construction. We discuss its limitations, and the obstacles facing a naive attempt of composing this graph with the long-code. In Section 3.3 we employ rather standard composition structure, and present a reduction from a gap-SAT instance

Φ to a graph $\mathcal{G}_{[\frac{1}{4},\frac{1}{8}]}(\Phi)$, showing NP-hardness for approximating Minimum-Vertex-Cover to within $\frac{7}{6}-\epsilon$ for any constant ϵ, the same factor whose hardness was already known [Hås97], although via a different construction. This construction exhibits the use of the p-biased long-code, which we believe is more clearly analyzed, but is a sidetrack and is not relied upon in the construction of the final graph $\mathcal{G}_{[p,p^\bullet]}(\Phi)$. Finally, in Section 3.4, we present a redesign of the manner by which the Composition technique [AS92] is applied. This new scheme starts with a phase in which Φ is preprocessed, coming up with a new set of variables that have a more *symmetric* structure, and whose consistency constraints are quite loose, nevertheless, one which can still be utilized by our analysis to deduce global consistency assuming local-consistency, which, in this case, translates to a small vertex cover. We apply the p-biased long-code over the new symmetric variables, to obtain $\mathcal{G}_{[p,p^\bullet]}(\Phi)$.

3.1 PCP Characterization of NP

PCP characterizations of NP in general state that given some SAT instance, namely a set of Boolean-functions $\Phi = \{\varphi_1, ..., \varphi_n\}$ over variables X, it is NP-hard to distinguish between the case where there is an assignment A to X that satisfies all $\varphi \in \Phi$, and the case where any assignment A satisfies at most a small fraction of Φ. The characterization used herein applies some specific requirements on the structure of Φ, described next.

Let $\langle \Phi, X, Y \rangle$ be a SAT instance over two types of variables, referred to as X and Y. Variables $x \in X$ take values in the range R_X while variables $y \in Y$ take values in R_Y. $\Phi = \{\varphi_1, ..., \varphi_n\}$ is a set of constraints (that is, Boolean-functions) over the variables X and Y, where each $\varphi \in \Phi$ is a constraint over one variable $x \in X$ and one variable $y \in Y$, $\varphi : R_X \times R_Y \to \{\mathsf{T}, \mathsf{F}\}$.

We say that Φ is *two-determined*, if for every $\varphi(x, y) \in \Phi$ and each value $a \in R_X$, there is exactly one value $b \in R_Y$ such that φ evaluates to True on $x = a$ and $y = b$.

For such x, y, we say that x *determines* y. Consequently, from now on, we would write $\varphi(x, y)$ as

$$\varphi_{x\to y} \colon R_X \to R_Y$$

We say that Φ is *regular*, if there exists an integer d_X so that each x determines exactly d_X y's, i.e. all $x \in X$ appear in exactly d_X constraints $\varphi_{x\to y} \in \Phi$:

$$\forall x \in X, \quad |\{\varphi_{x\to y} \in \Phi\}| = d_X .$$

An assignment is a function, assigning a value in R_X to each variable $x \in X$, and a value in R_Y to each variable $y \in Y$. We therefore use the notation $A \colon (X \to R_X, Y \to R_Y)$ to denote such an assignment, where $A(x)$ is the value assigned to a variable $x \in X$ while $A(y)$ is the value assigned to y.

Definition 3.1 *Let us denote by $\Upsilon(\Phi)$ the maximum, over all assignments to Φ's variables A: $(X \to R_X, Y \to R_Y)$, of the fraction of $\varphi \in \Phi$ satisfied by A, namely*

$$\Upsilon(\Phi) = \max_A \Pr_{\varphi_{x \to y}} \left[\varphi_{x \to y}(A(x)) = A(y)\right]$$

The PCP characterization we use herein is as follows:

Theorem 3.1 *Let $\epsilon > 0$ be any arbitrary constant. Let Φ be regular and two-determined. It is NP-hard to distinguish between the following two cases:*

- $\Upsilon(\Phi) = 1$*, i.e. Φ is satisfiable.*
- $\Upsilon(\Phi) < \epsilon$*, i.e. no assignment A satisfies a non-negligible fraction of Φ.*

Furthermore, this is true even when Φ is so that $|R_Y|, |R_X| \leq \epsilon^{-O(1)}$.

Proof: We prove this theorem by applying the parallel repetition lemma of [Raz98] to the gap-SAT instance of [AS92, ALM$^+$92]. The basic PCP theorem showing hardness for gap-SAT states that,

Theorem 3.2 ([AS92, ALM$^+$92]) *There exists some constant $\beta > 0$ such that given a set $\Psi = \{\psi_1, .., \psi_n\}$ of 3-CNF clauses over Boolean variables W (each clause is the OR of exactly 3 variables), it is NP-hard to distinguish between the two cases:*

- $\Upsilon(\Psi) = 1$
- $\Upsilon(\Psi) < 1 - \beta$

■

Let us define the parallel repetition version of Ψ,

Definition 3.2 ($\mathsf{Par}[\Psi, k]$) *Let $\langle \Psi, W \rangle$ be a 3-CNF instance, with 3-CNF clauses Ψ over variables W. For any integer $k > 0$, let*

$$\mathsf{Par}[\Psi, k] \overset{def}{=} \langle \Phi, X, Y \rangle$$

be a SAT instance with Boolean functions Φ over two types of variables: $X \overset{def}{=} \Psi^k$ and $Y \overset{def}{=} W^k$.

The range of each variable $x \in X$, is $R_X = [7]^k$, corresponding (by enumerating the 7 satisfying assignments of each 3-CNF clause $\psi \in \Psi$) to the concatenation of the satisfying assignments for Ψ's clauses in x. The range of each variable $y \in Y$, is $R_Y = [2]^k$, corresponding to all possible assignments to W's variables in y.

For $y = (w_1, .., w_k)$ and $x = (\psi_1, .., \psi_k)$, denote $y \sqsubseteq x$ if for all $i \in [k]$, w_i is a variable in ψ_i. The Boolean-functions in Φ are as follows:

$$\Phi = \left\{ \varphi_{x \to y} \;\middle|\; y \in W^k,\ x \in \Psi^k,\ y \sqsubseteq x \right\}$$

where $\varphi_{x \to y}$ is T if the assignment to y is the restriction to y of the assignment to x.

Note that Φ is regular, since every $x = (\psi_1, .., \psi_k)$ appears in $d_X = 3^k$ Boolean functions, one for each $y \sqsubseteq x$. Moreover, as hinted by the notation $\varphi_{x \to y}$, a value for x determines the one value for $y \sqsubseteq x$ that satisfies $\varphi_{x \to y} \in \Phi$.
Clearly, if $\Upsilon(\Psi) = 1$, then $\Upsilon(\Phi) = 1$. Moreover,

Theorem 3.3 (Parallel Repetition, [Raz98]) *For every $\beta > 0$ there exists some constant $c > 0$, such that the following holds. Let $\langle \Psi, W \rangle$ be a SAT-instance with $\Upsilon(\Psi) \leq 1 - \beta$, and for any $k > 0$ let $\langle \Phi, X, Y \rangle = \mathsf{Par}[\Psi, k]$. Then,*

$$\Upsilon(\Phi) \leq \Upsilon(\Psi)^{c \cdot k} \leq (1 - \beta)^{c \cdot k}$$

■

Taking k to be the first for which $(1 - \beta)^{c \cdot k} \leq \epsilon$ we indeed have $|R_Y| \leq |R_X| = \epsilon^{-O(1)}$. ■

For our purposes, this characterization needs to be slightly enhanced. The enhancement involves the notion of a multi-assignment - assigning a small set of values to Φ's variables - and what it means for such an assignment to "semi-satisfy" Φ's Boolean functions.

For any finite set R, denote the family of all size-h subsets of R by $\binom{R}{h} \overset{def}{=} \{F \subset R \mid |F| = h\}$.
An *h-assignment* A assigns to each variable in X and in Y a subset of h plausible values A: $(X \to \binom{R_X}{h}, Y \to \binom{R_Y}{h})$.

We say $\varphi_{x \to y}$ is *semi-satisfied* by an h-assignment A if there exists $a \in A(x)$ such that $\varphi_{x \to y}(a) \in A(y)$.

Let now Υ_h denote the maximum, over all h-assignments A, of the fraction of $\varphi \in \Phi$ semi-satisfied by A:

$$\Upsilon_h(\Phi) = \max_A \Pr_{\varphi_{x \to y} \in \Phi} \left[\varphi_{x \to y}(A(x)) \cap A(y) \neq \phi \right]$$

The PCP characterization of NP above (Theorem 3.1) can be extended to read as follows:

Corollary 3.4 *For any constants $h, \epsilon > 0$, given a SAT instance Φ that is regular and two-determined, it is NP-hard to distinguish between the following two cases:*

- $\Upsilon(\Phi) = 1$
- $\Upsilon_h(\Phi) < \epsilon$

Furthermore, this is true even when Φ *is so that* $|R_X|, |R_Y| \leq (\frac{h}{\epsilon})^{O(1)}$.

Proof: Let Φ be as in corollary 3.1, so that it is NP-hard to distinguish between the case where $\Upsilon(\Phi) = 1$ and the case where $\Upsilon(\Phi) < \epsilon \cdot \frac{1}{h^2}$. If $\Upsilon(\Phi) = 1$ then $\Upsilon_h(\Phi) = 1$. Suppose $\Upsilon_h(\Phi) > \epsilon$ and let A be an h-assignment semi-satisfying ϵ of Φ. Consider the assignments obtained by choosing randomly one, out of the h, value for each variable. The expected fraction of local-constraints such a random assignment satisfies is h^2 smaller than the fraction A semi-satisfies, and since at least one assignment must meet the expectation, $\Upsilon(\Phi) \geq \Upsilon_h(\Phi) \cdot \frac{1}{h^2} > \epsilon \cdot \frac{1}{h^2}$. ■

3.2 The FGLSS Graph

In this section we present a first attempt to prove the hardness of approximating Minimum-Vertex-Cover.

Let $\langle \Phi, X, Y \rangle$ be a gap-SAT instance as in Corollary 3.4, i.e. with Boolean constraints Φ over variables X and Y whose ranges are R_X and R_Y, such that either $\Upsilon(\Phi) = 1$ or $\Upsilon_h(\Phi) \leq \varepsilon_\Phi$. Recall that $\Upsilon(\Phi) = 1$ implies there is an assignment $A_\Phi \colon (X \to R_X; Y \to R_Y)$ satisfying all $\varphi_{x \to y}$ in Φ.

Let us consider the natural attempt, which is a variation on the FGLSS reduction [FGL+91, Kar72], for constructing a graph whose independent set is either large or small depending on which of the two cases Φ is in: let $\mathcal{G}_Z[\Phi]$ be the graph

$$\mathcal{G}_Z[\Phi] = \langle Z, E_Z \rangle \text{ where } Z \overset{def}{=} (X \times R_X)$$

that is, where $\mathcal{G}_Z[\Phi]$'s vertices is the set of pairs consisting of a variable x in X and a value $a \in R_X$ for x.

For the edge set E_Z of $\mathcal{G}_Z[\Phi]$, let us consider all pairs of vertices whose values cannot possibly correspond to the same satisfying assignment A_Φ. Let $\mathcal{I}[A_\Phi] \subseteq Z$ be the derivative of A_Φ, defined to be

$$\mathcal{I}[A_\Phi] = \{ (x, A_\Phi(x)) \mid x \in X \}$$

There are two types of inconsistency between vertices of Z to consider. Two distinct vertices (x, a_1) and (x, a_2), where x is common to both, cannot be both in $\mathcal{I}[A_\Phi]$ as A_Φ assigns a single value to x. Furthermore, if for some y, both x_1 and x_2 determine y, and assigning a_1 to x_1 fixes y to a value different to what is determined by x_2 assigned a_2; these two vertices cannot both be in $\mathcal{I}[A_\Phi]$. It turns out though, that the first inconsistency is a special case of the second, as two distinct values to x imply two distinct values for at least one y determined by x (assuming two values to x must differ on the value fixed to at least one y):

Definition 3.3 (Z-inconsistency) *Two vertices in Z, $z_1 = (x_1, a_1)$ and $z_2 = (x_2, a_2)$ are* Z-inconsistent*, if there is some variable $y \in Y$ with $\varphi_{x_1 \to y}, \varphi_{x_2 \to y} \in \Phi$ such that $\varphi_{x_1 \to y}(a_1) \neq \varphi_{x_2 \to y}(a_2)$. Let*

$$E_Z \overset{def}{=} \{\{z_1, z_2\} \mid z_1, z_2 \text{ are } Z\text{-inconsistent}\}.$$

Therefore, an independent set cannot correspond to an inconsistent assignment to Φ.

Note however that the gap between the sizes of the independent set, in case Φ is satisfiable and in case Φ is extremely not satisfiable, would not be good enough for our purposes. The size of the independent set in the first case would be $|X| = \frac{1}{|R_X|} \cdot |Z|$ compared to ϵ_Φ times that in the second case, translating to a ratio of $\frac{1-\epsilon/|R_X|}{1-1/|R_X|}$ for vertex-cover, which is significantly smaller than our goal.

Here might be the place to note that Håstad's obtained his hardness result for Minimum-Vertex-Cover by taking the FGLSS graph over his Linear-Equations gap-SAT instance, which is equivalent (ignoring the imperfect completeness that is irrelevant here) to $|R_X| = 4$ and $\epsilon = \frac{1}{2}$, yielding a hardness factor of $\frac{1-\frac{1}{2}\cdot\frac{1}{4}}{1-\frac{1}{4}} = \frac{7}{6}$.

A Naive Construction. Let us first consider a naive manner by which to apply the long-code to the graph $\mathcal{G}_Z[\Phi]$. Denote the set of vertices of $\mathcal{G}_Z[\Phi]$ that correspond to a variable $x \in X$ by:

$$Z[x] \overset{def}{=} \{(x, a) \in Z \mid a \in R_X\} = \{x\} \times R_X$$

Since an independent set in $\mathcal{G}_Z[\Phi]$ has at most one representative in $Z[x]$ for each $x \in X$, and the maximal independent set contains exactly one representative in each $Z[x]$, we could replace each set of vertices $Z[x]$ with a set of vertices corresponding to the long-code of $Z[x]$, supposedly encoding $Z[x]$'s representative in $\mathcal{G}_Z[\Phi]$'s maximal independent set. Edges in this graph connect two subsets that cannot possibly be consistent, namely, $F_1 \in \mathbf{P}(Z[x_1])$ is connected to $F_2 \in \mathbf{P}(Z[x_2])$ iff $F_1 \cup F_2$ form a clique in $\mathcal{G}_Z[\Phi]$, in other words, if

$$\forall a_1 \in F_1, a_2 \in F_2, \qquad \exists y, \varphi_{x_1 \to y}, \varphi_{x_2 \to y} \in \Phi, \ s.t.\ \varphi_{x_1 \to y}(a_1) \neq \varphi_{x_2 \to y}(a_2)$$

This ensures that the maximal independent set in this graph would be of large size, in case Φ is satisfiable. However, it might also be of large size in case Φ is far from satisfiable as well. To see this, consider for candidates in the independent set, for each $x \in X$, only subsets $F \in \mathbf{P}(R_X)$ for which $\varphi_{x \to y}(F) = R_Y$ for every $\varphi_{x \to y} \in \Phi$. Such subsets contain, for every y determined by x and every value $b \in R_Y$ for y, at least one element $a \in F$ with $\varphi_{x \to y}(a) = b$. On the one hand, any two such subsets, for two distinct x's, are consistent. On the other hand, almost all subsets satisfy this requirement. To see this, note that the probability that a subset $F \in \mathbf{P}(R_X)$ misses all of $\varphi^{-1}_{x \to y}(b)$ is exponentially small in the number of a's mapped to one b; which is negligible even if multiplying by the number of y's determined by x. Hence one can construct a large independent set corresponding to an arbitrary assignment a_x to each x, by taking only those subsets that contain a_x and, in addition, satisfy the above requirement.

3.3 A 7/6 Construction

The naive construction of taking the long-code of each variable in X, failed to enforce consistency between the encodings of two such x's. One possible way to go around this, is by gluing together the encodings of both variables x and y for each $\varphi_{x\to y} \in \Phi$, in an adaptive manner defined below. Consistency is then easily verified by inserting edges between inconsistent encodings of each appearance of x or of y. This construction achieves consistency between tests, yet loses in that the size of the independent in case Φ is satisfiable is $\frac{1}{4}$ rather than $\frac{1}{2}$, compared to $\frac{1}{8}$ in case Φ is far from satisfiable. This achieves a hardness of approximation ratio of roughly $\frac{1-\frac{1}{8}}{1-\frac{1}{4}} = \frac{7}{6}$ for vertex-cover. Before we continue, let us note that Hastad [Hås97] proved that vertex-cover is hard to approximate to within $\frac{7}{6}$, by constructing the FGLSS graph over the Linear-Equations system. The construction we present here is an alternative to that result.

Again, let $\langle \Phi, X, Y\rangle$ be regular and two-determined as in Corollary 3.4, and let $\varphi_{x\to y} \in \Phi$. An assignment to $\varphi_{x\to y}$ consists of a value in R_X and a value in R_Y, hence the long code of $\varphi_{x\to y}$ would be $\mathbf{P}(R_X) \times \mathbf{P}(R_Y) \cong \mathbf{P}(R_Y \cup R_X)$. We enforce consistency between x and y by allowing only those subsets $F \in \mathbf{P}(R_Y \cup R_X)$ for which, denoting $F_X = F \cap R_X$ and $F_Y = F \cap R_Y$,

$$\varphi_{x\to y}(F_X) \subseteq F_Y$$

We assign these sets the distribution obtained by first selecting a random subset $F_Y \in_{\mu_p} \mathbf{P}(R_Y)$ and then selecting a random subset $F_X \in_{\mu_p} \mathbf{P}\left(\varphi^{-1}_{x\to y}(F_Y)\right)$.

Definition 3.4 *The p-adaptive mutual long-code of $\varphi_{x\to y}$ is the family of subsets*

$$\mathcal{LC}^{x\to y} \overset{def}{=} \{F \in \mathbf{P}(R_X \cup R_Y) \mid F_Y \supseteq \varphi_{x\to y}(F_X)\}$$

endowed with the distribution

$$\mu_p^{x\to y}(F) \overset{def}{=} \mu_p^{R_Y}(F_Y) \cdot \mu_p^{\varphi^{-1}_{x\to y}(F_Y)}(F_X)\,.$$

Define the graph

$$\mathcal{G}_{[\frac{1}{4},\frac{1}{8}]}(\Phi) \overset{def}{=} \left\langle V_{[\frac{1}{4},\frac{1}{8}]}, E_{[\frac{1}{4},\frac{1}{8}]}\right\rangle$$

with vertices

$$V_{[\frac{1}{4},\frac{1}{8}]} = \bigcup_{\varphi_{x\to y}\in\Phi} \mathcal{LC}^{x\to y}$$

(we take $\mathcal{LC}^{x\to y}$ to be distinct for distinct tests $\varphi_{x\to y} \in \Phi$), whose weights are defined to be, setting $p = \frac{1}{2} - \epsilon$,

$$\forall F \in \mathcal{LC}^{x\to y}, \qquad \Lambda(F) \overset{def}{=} \frac{1}{|\Phi|} \cdot \mu_p^{x\to y}(F)$$

A pair of subsets $F \in \mathcal{LC}^{x\to y}$ and $F' \in \mathcal{LC}^{x\to y'}$ are connected by an edge if $F \cap F' \cap R_X = \phi$ (this applies also to the case where $y = y'$),

$$E_1 \overset{def}{=} \left\{ (F, F') \in \mathcal{LC}^{x\to y} \times \mathcal{LC}^{x\to y'} \;\middle|\; F \cap F' \cap R_X = \phi \right\} .$$

For distinct variables $x_1, x_2 \in X$ for which $\varphi_{x_1\to y}, \varphi_{x_2\to y} \in \Phi$, a pair of vertices $F_1 \in \mathcal{LC}^{x_1\to y}, F_2 \in \mathcal{LC}^{x_2\to y}$ are connected by an edge if $\varphi_{x_1\to y}(R_X \cap F_1) \cap \varphi_{x_2\to y}(R_X \cap F_2) = \phi$,

$$E_2 \overset{def}{=} \{ (F_1, F_2) \in \mathcal{LC}^{x_1\to y} \times \mathcal{LC}^{x_2\to y} \mid \varphi_{x_1\to y}(R_X \cap F_1) \cap \varphi_{x_2\to y}(R_X \cap F_2) = \phi \} .$$

Altogether,

$$E_{[\frac{1}{4},\frac{1}{8}]} \overset{def}{=} E_1 \cup E_2 .$$

Now, in case $A_\Phi \colon (X \to R_X; Y \to R_Y)$ is a satisfying assignment for Φ, taking for each $\varphi_{x\to y}$ the family of all subsets containing both $A(x), A(y)$ comprises an independent set in $\mathcal{G}_{[\frac{1}{4},\frac{1}{8}]}(\Phi)$, of weight $p^2 > \frac{1}{4} - 2\epsilon$:

Lemma 3.5 (Completeness of $\mathcal{G}_{[\frac{1}{4},\frac{1}{8}]}(\Phi)$) *If $A_\Phi \colon (X \to R_X; Y \to R_Y)$ is a satisfying assignment for Φ, then the following set is an independent set in $\mathcal{G}_{[\frac{1}{4},\frac{1}{8}]}(\Phi)$,*

$$\mathcal{I}[A_\Phi] = \bigcup_{\varphi_{x\to y}\in\Phi} \{ F \in \mathcal{LC}^{x\to y} \mid A_\Phi(x), A_\Phi(y) \in F \}$$

Note the inherent loss from the composed structure of $\mathcal{LC}^{x\to y}$, halving the size of the independent set twice, once for x and once for y.

For the soundness of this construction, we show that if $\mathcal{G}_{[\frac{1}{4},\frac{1}{8}]}(\Phi)$ has an independent set whose weight is even slightly more than $\frac{1}{8}$, we can find an h-assignment semi-satisfying more than ϵ of Φ, thus Φ is satisfiable:

Lemma 3.6 (Soundness of $\mathcal{G}_{[\frac{1}{4},\frac{1}{8}]}(\Phi)$) *For any SAT instance Φ, if $\alpha(\mathcal{G}_{[\frac{1}{4},\frac{1}{8}]}(\Phi)) \geq \frac{1}{8} + 8\epsilon$, then $\Upsilon_h(\Phi) \geq \epsilon$ where $h = \mathsf{h}(\frac{1}{2} - \epsilon, \epsilon, \frac{4}{\epsilon}) = O(1)$.*

Proof: Let $\mathcal{I} \subset V_{[\frac{1}{4},\frac{1}{8}]}$ be an independent set whose weight is $\Lambda(\mathcal{I}) \geq \frac{1}{8} + 8\epsilon$. The proof proceeds in three steps:

In the first step, we observe that for every x_0, y_0 the families

$$\mathcal{I}_{x_0} = \bigcup_{y\colon \varphi_{x_0\to y}\in\Phi} \{ F \cap R_X \mid F \in \mathcal{I} \cap \mathcal{LC}^{x_0\to y} \} \quad \text{and} \quad \mathcal{I}_{y_0} = \bigcup_{x\colon \varphi_{x\to y_0}\in\Phi} \{ F \cap R_Y \mid F \in \mathcal{I} \cap \mathcal{LC}^{x\to y_0} \}$$

are intersecting because of the edges in E_1 and E_2 respectively. Moreover, assuming w.l.o.g. that $\mathcal{I}$ is maximal, these families are also monotone.

The second step would be to find a significant portion of the tests in Φ, and for each of their variables, a small set of permissible values. Let $\Phi_1 \subseteq \Phi$ consist of all $\varphi_{x\to y} \in \Phi$ for which $\mu_p(\mathcal{I} \cap \mathcal{LC}^{x\to y}) > \frac{1}{8} + 6\epsilon$, noting $|\Phi_1| \geq 2\epsilon \cdot |\Phi|$.

Next, we find a subset $\Phi_\epsilon \subseteq \Phi_1$ of significant size, $|\Phi_\epsilon| \geq \epsilon \cdot |\Phi|$, and some $q \in (p, p+\epsilon)$ such that for every $\varphi_{x\to y} \in \Phi_\epsilon$, $\mathbf{as}_q(\mathcal{I}_y) + \mathbf{as}_q(\mathcal{I}_x) < \frac{4}{\epsilon}$. Such a set must exist since

$$\frac{1}{|\Phi_1|} \sum_{\varphi_{x\to y} \in \Phi_1} \mathbf{as}_p(\mathcal{I}_x) + \mathbf{as}_p(\mathcal{I}_y) = \frac{d}{dp}\left(\frac{1}{|\Phi_1|} \sum_{\varphi_{x\to y} \in \Phi_1} \mu_p(\mathcal{I}_x) + \mu_p(\mathcal{I}_y)\right)$$

and the right hand side, being the derivative of a function that is bounded between 0 and 2, must be smaller than $\frac{2}{\epsilon}$ at some point $q \in (p, p+\epsilon)$. At most half of the tests $\varphi_{x\to y} \in \Phi_1$ can have $\mathbf{as}_q(\mathcal{I}_y) + \mathbf{as}_q(\mathcal{I}_x) > 2 \cdot \frac{2}{\epsilon}$ which is twice the expectation, thus $|\Phi_\epsilon| \geq \frac{1}{2} \cdot |\Phi_1| \geq \epsilon\,|\Phi|$.

Denote by X_ϵ (resp. Y_ϵ) the X-variables (resp. Y-variables) appearing in the tests of Φ_ϵ. For every variable $z \in X_\epsilon \cup Y_\epsilon$, we now apply the Friedgut-Lemma (Theorem 2.3) to find a constant sized (ϵ, q)-core for $\mathcal{I}_z$, C_z, $|C_z| \leq h = \mathsf{h}(\frac{1}{2} - \epsilon, \epsilon, \frac{4}{\epsilon}) = O(1)$, and define an h-assignment for Φ by setting $A(z) \overset{def}{=} C_z$ for each $z \in X_\epsilon \cup Y_\epsilon$ and for the remaining variables, simply $A(z) \overset{def}{=} \phi$.

The third and final step is to prove that this h-assignment semi-satisfies Φ. This is done by showing that, for every y, x with $\varphi_{x\to y} \in \Phi_\epsilon$, $\varphi_{x\to y}(A(x)) \cap A(y) \neq \phi$. So assume otherwise and recall from Definition 2.5 the core families of $\mathcal{I}_x$ and $\mathcal{I}_y$, denoted $\langle \mathcal{I}_x \rangle_{C_x}$ and $\langle \mathcal{I}_y \rangle_{C_y}$ respectively, and define

$$CF_x = \left\{ F_X \subset R_X \mid F_X \cap C_x \in \langle \mathcal{I}_x \rangle_{C_x} \right\} \quad \text{and} \quad CF_y = \left\{ F_Y \subset R_Y \mid F_Y \cap C_y \in \langle \mathcal{I}_y \rangle_{C_y} \right\}.$$

By Definition 3.4, and by Lemma 2.1,

$$\begin{aligned} \mu_p^{x\to y}(\mathcal{I} \cap \mathcal{LC}^{x\to y}) \;&=\; \sum_{F \in \mathcal{I} \cap \mathcal{LC}^{x\to y}} \mu_p^{\varphi^{-1}_{x\to y}(F \cap R_Y)}(F \cap R_X) \cdot \mu_p^{R_Y}(F \cap R_Y) \\ &\leq\; 2 \cdot 3\epsilon + \sum_{\substack{F \,\in\, \mathcal{I} \,\cap\, \mathcal{LC}^{x\to y} \\ F_X \,=\, F \,\cap\, R_X \,\in\, CF_x \\ F_Y \,=\, F \,\cap\, R_Y \,\in\, CF_y}} \mu_p^{\varphi^{-1}_{x\to y}(F_Y)}(F_X) \cdot \mu_p^{R_Y}(F_Y)\,. \end{aligned}$$

For every fixed $F_Y \in CF_y$, consider the family

$$\mathcal{F}[F_Y] \overset{def}{=} \left\{ F_X \subset \varphi^{-1}_{x\to y}(F_Y) \;\middle|\; F_X \in CF_x,\ F_Y \cup F_X \in \mathcal{I} \cap \mathcal{LC}^{x\to y} \right\},$$

and observe that it is intersecting, therefore $\mu_p^{\varphi^{-1}_{x\to y}(F_Y)}(\mathcal{F}[F_Y]) \leq p$. Moreover, the above sum is equal to

$$= 6\epsilon + \sum_{F_Y \in CF_y} \mu_p^{R_Y}(F_Y) \cdot \sum_{F_X \in \mathcal{F}[F_Y]} \mu_p^{\varphi^{-1}_{x\to y}(F_Y)}(F_X) \leq 6\epsilon + \sum_{F_Y \in CF_y,\, \mathcal{F}[F_Y] \neq \phi} \mu_p^{R_Y}(F_Y) \cdot p$$

We complete the proof by showing that $\mu_p^{R_Y}(\{ F_Y \in CF_y \mid \mathcal{F}[F_Y] \neq \phi \}) \leq p^2 < \frac{1}{4}$ contradicting the fact that for any $\varphi_{x\to y} \in \Phi_\epsilon$, $\mu_p(\mathcal{I} \cap \mathcal{LC}^{x\to y}) > \frac{1}{8} + 6\epsilon$. Let $\varphi_{x\to y}(CF_x) = \{ \varphi_{x\to y}(F_X) \mid F_X \in CF_x \}$, and observe that this family is monotone and intersecting and is completely determined (see Section 2.1) by $\varphi_{x\to y}(C_x)$. Finally, if $\mathcal{F}[F_Y] \neq \phi$ then $F_Y \in \varphi_{x\to y}(CF_x)$ because $\exists F_X \in \mathcal{F}[F_Y] \subset CF_x$ which means $\varphi_{x\to y}(F_X) \subseteq F_Y$ and by monotonicity. Thus, the above comes to

$$\mu_p(\{ F_Y \in CF_y \mid \mathcal{F}[F_Y] \neq \phi \}) \leq \mu_p(\varphi_{x\to y}(CF_x) \cap CF_y) = \mu_p(\varphi_{x\to y}(CF_x)) \cdot \mu_p(CF_y) \leq p^2$$

where the equality holds due to $\varphi_{x\to y}(C_x) \cap C_y = \varphi_{x\to y}(A(x)) \cap A(y) = \phi$ and because μ_p is a product measure; and the last inequality comes from each of CF_y and $\varphi_{x\to y}(CF_x)$ being an intersecting family which limits their size to p.

We have reached a contradiction, proving that for every $\varphi_{x\to y} \in \Phi_\epsilon$, $\varphi_{x\to y}(A(x)) \cap A(y) \neq \phi$, and so $\Upsilon_h(\Phi) \geq \epsilon$. ■

3.4 The Final Graph, $\mathcal{G}_{[p,p^\bullet]}(\Phi)$

The construction of $\mathcal{G}_{[\frac{1}{4},\frac{1}{8}]}(\Phi)$ is inherently limited by the fact that we glued the long-codes of x and of y together, twice halving the size of the largest possible independent set, even in case Φ is satisfiable. It seems that achieving hardness for a factor of approximation larger than $\frac{7}{6}$, and especially reaching $2 - \epsilon$, requires encoding each variable separately.

Recall however, the naive attempt (described in Section 3.2), of separately taking the long-code over the possible values R_X for each $x \in X$, and then adding consistency edges between distinct x's. This attempt could not enforce consistency between the long-codes of distinct x's due to the following. Consider a pair of variables x_1, x_2 that both determine y. $\varphi_{x_1\to y}$ (and similarly $\varphi_{x_2\to y}$) partitions the values $a \in R_X$ according to $\varphi_{x_1\to y}(a)$, such that for every $b \in R_Y$, every value $a_1 \in \varphi_{x_1\to y}^{-1}(b)$ is consistent with every value $a_2 \in \varphi_{x_2\to y}^{-1}(b)$. The fake consistency stems from the fact that every value $b \in R_Y$ has many possible origins in R_X, each showing up in $F \in_{\mu_p} \mathbf{P}(R_X)$ independently with probability p. Thus a typical subset $F \in \mathcal{LC}_p^{R_X}$, when projected to R_Y, usually covers all possible values in R_Y, i.e.

$$\Pr_{F \in \mathcal{LC}_p^{R_X}} [\varphi_{x\to y}(F) = R_Y] \approx 1$$

Such subsets are consistent with each other.

Our next attempt, is to construct new variables $\mathcal{Z}$, such that setting a value for one variable leaves at most two possible values for related variables. Together with the specific structure of $\mathcal{Z}$, we would be able, taking p to be small enough $p < \frac{3-\sqrt{5}}{2}$, to establish consistency between p-biased long-codes of different variables.

The constructed graph $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ will be structured according to the FGLSS graph, $\mathcal{G}_Z[\Phi] = \langle Z, E_Z \rangle$, as constructed in Section 3.2; and will implicitly depend on a parameter

l to be fixed later (see Definition 4.1) so that $l \geq 8 \log \frac{2}{\epsilon} \cdot |R_X|$. Consider the family $\mathcal{Z}$ of all sets of size l of Z:

$$\mathcal{Z} = \binom{Z}{l}$$

Let us refer to each $B \in \mathcal{Z}$ as a *block*. Recall we denoted by $\mathcal{I}[A_\Phi] \overset{def}{=} \{(x, A_\Phi(x))\}$ the set of vertices in $\mathcal{G}_Z[\Phi]$ corresponding to an assignment A_Φ for Φ. The intersection of $\mathcal{I}[A_\Phi]$ with any $B \in \mathcal{Z}$, $\mathcal{I}[A_\Phi] \cap B$, can take 2^l distinct forms, namely all subsets of B. Consider then an assignment $A_{\mathcal{Z}}$ to the set of blocks $B \in \mathcal{Z}$, in which each block B is assigned a truth assignment

$$A_{\mathcal{Z}}(B) \in \{\mathsf{f} \colon B \to \{\mathsf{T}, \mathsf{F}\}\}$$

supposedly assigning T to exactly all vertices of B that are in the independent set $\mathcal{I}[A_\Phi]$, that is, where $\mathsf{f}^{-1}(\mathsf{T}) = \mathcal{I}[A_\Phi] \cap B$. For a truth assignment for B, $\mathsf{f} \colon B \to \{\mathsf{T}, \mathsf{F}\}$, and any $\hat{B} \subseteq B$, let us denote by $\mathsf{f}|_{\hat{B}} \colon \hat{B} \to \{\mathsf{T}, \mathsf{F}\}$ the restriction of f to $\hat{B}$, namely, where $\forall z \in \hat{B}$, $\mathsf{f}|_{\hat{B}}(z) = \mathsf{f}(z)$. Given a pair of blocks B_1, B_2 that intersect on $\hat{B} = B_1 \cap B_2$ with $|\hat{B}| = l - 1$, every truth-assignment to B_1 is consistent with exactly *two* truth-assignments to B_2.

Now, with little more effort, we can ensure the truth-assignments assigned to the blocks in $\mathcal{Z}$ correspond to a large independent set; this follows from the following observation regarding the restriction of $\mathcal{I}[A_\Phi]$ to all $B \in \mathcal{Z}$. Since $|\mathcal{I}[A_\Phi]| = |X| = \frac{|Z|}{|R_X|}$, the expected number of vertices assigned T in the restriction of $\mathcal{I}[A_\Phi]$ to a random $B \in \mathcal{Z}$, is $\frac{l}{|R_X|}$. Let $t = \frac{1}{2} \frac{l}{|R_X|}$ be half of this expectation, and consider the probability that B has fewer than t vertices in $\mathcal{I}[A_\Phi]$; by a straightforward application of the Chernoff bound (see Proposition B.2 in the appendix), for any $\mathcal{I}[A_\Phi]$

$$\Pr_{B \in \mathcal{Z}} \left[|\mathcal{I}[A_\Phi] \cap B| < t \right] < 2e^{-\frac{l}{8|R_X|}}$$

It is therefore the case that, disallowing $A_{\mathcal{Z}}$ to assign blocks $B \in \mathcal{Z}$ a truth-assignment over B, if that assignment assigns T to fewer than t vertices $z \in B$, excludes only a tiny fraction of the blocks from being assigned the truth-assignment that corresponds to $\mathcal{I}[A_\Phi] \cap B$, and allows $A_{\mathcal{Z}}$ to be consistent almost everywhere. Consequently, let, for each block B,

$$R_B \overset{def}{=} \left\{ \mathsf{f} \colon B \to \{\mathsf{T}, \mathsf{F}\} \;\middle|\; |\mathsf{f}^{-1}(\mathsf{T})| \geq t \right\}$$

and consider $A_{\mathcal{Z}}$ that assigns to each $B \in \mathcal{Z}$ only values in R_B, $A_{\mathcal{Z}}(B) \in R_B$. As proved below, this restriction ensures a sizeable set in $\mathcal{G}_Z[\Phi]$, and, if consistent, being an independent set, must correspond to an assignment A satisfying sizeable fraction of Φ. Let us refer, from now on, to each member of R_B as a *block assignment*.

So the next attempt would be to apply the FGLSS reduction to the R_B's, and construct a graph $\mathcal{G}_{\mathcal{Z}}[\Phi]$ whose vertices consist of all block-assignments to every block $B \in \mathcal{Z}$, and whose edges correspond to pairs of such block assignments that cannot possibly be consistent with an independent-set in $\mathcal{G}_Z[\Phi]$:

$$\mathcal{G}_{\mathcal{Z}}[\Phi] = \left\langle \bigcup_{B \in \mathcal{Z}} R_B \, , \, \bowtie \right\rangle$$

where the edge relation $\bowtie$ is defined by stating when two block-assignments are considered inconsistent:

Definition 3.5 (Inconsistent Block-assignments) *Let $B_1, B_2 \in \mathcal{Z}$ and denote $\hat{B} = B_1 \cap B_2$; a pair of block-assignments $\mathsf{f}_1 \in R_{B_1}, \mathsf{f}_2 \in R_{B_2}$ are* inconsistent*, denoted $\mathsf{f}_1 \bowtie \mathsf{f}_2$, if either $\mathsf{f}_1|_{\hat{B}} \neq \mathsf{f}_2|_{\hat{B}}$, or, if there exist $z_1 \in B_1$ and $z_2 \in B_2$ that are Z-inconsistent, and such that $\mathsf{f}_1(z_1) = \mathsf{f}_2(z_2) = \mathsf{T}$.*

The composition technique is now called upon to encode each block-assignment of every $B \in \mathcal{Z}$, applying the biased–long-code to each, and check their consistency. This is still not an easy task, since the consistency between distinct B's can be checked only for blocks that differ only on one element of Z, but these hardships belong to a different section.

The Constructed Graph

It is now time to define our graph $\mathcal{G}_{[p,p^{\bullet}]}(\Phi)$.

Vertices and Weights: $\mathcal{G}_{[p,p^{\bullet}]}(\Phi) = \langle V, E, \Lambda \rangle$ has a block of vertices $V[B]$ for every $B \in \mathcal{Z}$, where vertices in each block B correspond to the p-biased–long-code applied to R_B

$$V[B] = \mathbf{P}\left(R_B\right)$$

that is, one vertex for each subset $F \subseteq R_B$ of B's block-assignments. V consists of one such block of vertices for each $B \in \mathcal{Z}$

$$V = \bigcup_{B \in \mathcal{Z}} V[B]$$

Note that we take the block-assignments to be distinct, hence, subsets of them are distinct, and V is a disjoint union of $V[B]$ over all $B \in \mathcal{Z}$.

Let Λ_B, for each block $B \in \mathcal{Z}$, be the distribution assigning each vertex F, a probability according to μ_p, namely

$$\Lambda_B(F) = \mu_p^{R_B}(F)$$

The block of vertices $V[B]$ superimposed with Λ_B therefore comprise a p-biased–long-code over R_B (see Definition 2.3).

The probability distribution Λ assigns uniform probability to each block: For any $F \in V[B]$

$$\Lambda(F) \overset{def}{=} |\mathcal{Z}|^{-1} \cdot \Lambda_B(F)$$

Edges. Within a block B any two non-intersecting vertices are connected by an edge:

$$E[B] \overset{def}{=} \left\{ (F_1, F_2) \in V[B]^2 \,\middle|\, F_1 \cap F_2 = \phi \right\}$$

A pair of vertices in distinct blocks, $F_1 \in V[B_1]$ and $F_2 \in V[B_2]$, are connected by an edge if all their block-assignments are pairwise inconsistent:

$$E[B_1, B_2] \overset{def}{=} \{(F_1, F_2) \in V[B_1] \times V[B_2] \ \mid \ \forall \mathsf{f}_1 \in F_1, \mathsf{f}_2 \in F_2 \quad \mathsf{f}_1 \bowtie \mathsf{f}_2 \}$$

Altogether,

$$E \overset{def}{=} \bigcup_{B \in \mathcal{Z}} E[B] \ \cup \bigcup_{B_1, B_2 \in \mathcal{Z}} E[B_1, B_2]$$

This completes the construction of the graph $\mathcal{G}_{[p,p^\bullet]}(\Phi)$.

Proposition 3.7 *The graph $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ is polynomial-time constructible given input Φ.* ■

Chapter 4

Main Theorem

The graph $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ constructed in the previous section, has an independent set of size almost p, in case Φ is satisfiable, as will be proven in Lemma 4.3 below. Furthermore, the heart of this chapter is to show that if Φ is far from being satisfiable, then $\alpha(\mathcal{G}_{[p,p^\bullet]}(\Phi))$ cannot be even slightly larger than $p^\bullet = \max(p^2, 4p^3 - 3p^4)$, provided that $p < p_{max} \overset{def}{=} \frac{3-\sqrt{5}}{2} \approx 0.382$. Thus,

Theorem 4.1 *Let $p < p_{max}$. For any constant $\varepsilon > 0$, given a weighted graph G, it is NP-hard to distinguish between the case where $\alpha(G) > p - \varepsilon$, and the case where $\alpha(G) < p^\bullet + \varepsilon$.*

Throughout the proof one may think of $p \leq \frac{1}{3}$ in which case $p^\bullet$ reads p^2. In this special case, the gap in our main theorem approaches $\frac{1-p^2}{1-p} = 1 + p$, yielding a hardness-of-approximation factor of $\frac{4}{3}$ for Minimum-Vertex-Cover. Before we prove the main theorem, let us state the corollary for Minimum-Vertex-Cover,

Corollary 4.2 (Minimum Vertex Cover) *Given a graph G, it is NP-hard to approximate $\overline{\alpha}(G)$ to within a factor of* 1.361.

Proof: For p near p_{max}, $p^\bullet = 4p^3 - 3p^4$, thus Theorem 4.1 asserts that it is NP-hard to distinguish between the case G has a vertex cover of size $1 - p + \epsilon$ and the case G has a vertex cover of size at least $1 - 4p^3 + 3p^4 - \epsilon$ for any $\epsilon > 0$. Minimum Vertex-Cover is thus shown hard to approximate to within a factor approaching

$$\frac{1 - 4(p_{max})^3 + 3(p_{max})^4}{1 - p_{max}} = 1 + p_{max} + (p_{max})^2 - 3(p_{max})^3 > 1.361$$

■

Proof: (of Theorem 4.1) The proof proceeds by reduction from the NP-complete problem gap-SAT, as in Corollary 3.4. Given $p, \varepsilon > 0$, we choose h, ε_Φ, and l (see Definition 4.1 below) and then, given Φ, construct the graph $\mathcal{G}_{[p,p^\bullet]}(\Phi)$. The completeness of the reduction, asserting that if $\Upsilon(\Phi) = 1$ then $\alpha(\mathcal{G}_{[p,p^\bullet]}(\Phi)) > p - \varepsilon$, is shown in Lemma 4.3. The soundness is established by Lemma 4.4, showing that if $\Upsilon_h(\Phi) < \varepsilon_\Phi$ then $\alpha(\mathcal{G}_{[p,p^\bullet]}(\Phi)) < p^\bullet + \varepsilon$. ■

We begin by setting the parameters. It is worthwhile to note here that the particular values chosen for these parameters are not important. They are chosen to satisfy some properties through the course of the proof, nevertheless, most importantly, aside from l, they are unrelated to $|R_X|$. Given ε and p, we can select values for h and ε_Φ that allow us to translate a large enough independent set into an h-assignment semi-satisfying ε_Φ of Φ. The size of R_X is affected by this choice, as $|R_X| = (\frac{\varepsilon_\Phi}{h})^{O(1)}$, therefore it is crucial that h depend only on ε and p, and *not* on $|R_X|$.

Definition 4.1 (Parameter Setting) *Let us set the parameters as follows.*

- *Let* $0 < \gamma < p_{max} - p$ *be such that,* $(p + \gamma)^\bullet - p^\bullet < \frac{1}{4}\varepsilon$.
- *We choose* h *to accommodate applications of Friedgut's Lemma, a Sunflower Lemma and a pigeon-hole principle. Recall that* $\mathsf{h}(p, \delta, k) \leq (c_p)^{k/\delta}$ *denoted the bound on the size of a* (δ, p)*-core of a family whose average sensitivity is bounded by* k*. Let*

$$h_0 = \sup_{q \in [p, p_{max}]} \left(\mathsf{h}(q, \tfrac{1}{16}\varepsilon, \tfrac{2}{\gamma}) \right)$$

  and let $\eta = \frac{1}{16h_0} \cdot p^{8h_0}$, $h_1 = \left\lceil \frac{2}{\gamma \cdot \eta} \right\rceil + h_0$, $h_s = 1 + 2^{2h_0} \cdot \sum_{k=0}^{h_0} \binom{h_1}{k}$, *and* $h = (h_s)^{h_1} \cdot (h_1)!$.
- *Fix* $\varepsilon_\Phi = \frac{1}{32h} \cdot \varepsilon$.
- *Fix* $l \overset{def}{=} \max(8 \log \frac{2}{\epsilon} \cdot |R_X|\,,\ 2(h_1)^2 \cdot |R_X|)$.

Remarks. The supremum $\sup_{q \in [p, p_{max}]} \left(\mathsf{h}(q, \frac{1}{16}\varepsilon, \frac{2}{p_{max} - p}) \right)$ in the definition of h_0 is bounded, because $\mathsf{h}(q, \frac{1}{16}\varepsilon, 16\varepsilon^{-1})$ is a continuous function of q, see Theorem 2.3. The value of γ is well defined because the function $f(p) = \max(p^2, 4p^3 - 3p^4)$ is a continuous function of p. Also, since the parameters affecting l, including $|R_X|$, (see Corollary 3.4) are constant and unrelated to the size of the instance $|\Phi|$, we can assume that $l < \frac{1}{2}|X|$.

4.1 Completeness

We now proceed to establish the (easier) completeness part of the reduction,

Lemma 4.3 (Completeness) *If* $\Upsilon(\Phi) = 1$ *then there exists an independent set* $\mathcal{I} \subset V$*, with*

$$\Lambda(\mathcal{I}) \geq p - \varepsilon$$

Proof: Let $A_\Phi\colon (X \to R_X, Y \to R_Y)$ be a satisfying assignment for Φ. Consider all vertices of $\mathcal{G}_Z[\Phi]$ consistent with A_Φ

$$\mathcal{I}[A_\Phi] = \{(x, A(x)) \in Z\}\,.$$

Let $\mathcal{Z}'$ be the set of blocks $B \in \mathcal{Z} = \binom{Z}{l}$ that intersect $\mathcal{I}[A_\Phi]$ on at least t elements $|B \cap \mathcal{I}[A_\Phi]| \geq t$. The probability that a block is not in $\mathcal{Z}'$ is bounded, via a simple Chernoff bound (see Proposition B.2), by $2e^{-\frac{l}{8|R_X|}} < \varepsilon$. For a block $B \in \mathcal{Z}'$, let us define the block-assignment $\mathsf{f}_B \in R_B$, assigning T to values from $\mathcal{I}[A_\Phi]$ and F to the rest:

$$\forall (x,a) \in B \qquad \mathsf{f}_B(x,a) \overset{def}{=} \begin{cases} \mathsf{T} & A_\Phi(x) = a \\ \\ \mathsf{F} & A_\Phi(x) \neq a \end{cases}$$

Now take $\mathcal{I}$ to be the set of all vertices containing f_B:

$$\mathcal{I} = \bigcup_{B \in \mathcal{Z}'} \{F \in V[B] \mid F \ni \mathsf{f}_B\}$$

To see that $\mathcal{I}$ is an independent set, first note that for every block $B \in \mathcal{Z}'$, $\mathcal{I} \cap V[B] \ni \mathsf{f}_B$. Furthermore, consider two blocks $B_1, B_2 \in \mathcal{Z}$ and a pair of vertices $F_1 \in V[B_1] \cap \mathcal{I}, F_2 \in V[B_2] \cap \mathcal{I}$ and denote $\hat{B} = B_1 \cap B_2$. $\mathsf{f}_{B_1}, \mathsf{f}_{B_2}$ must coincide on $\hat{B}$ as both are defined according to A. If $B_1 \ni (x_1, a_1)$ and $B_2 \ni (x_2, a_2)$ are Z-inconsistent (see definition 3.3), since A_Φ is a satisfying assignment, not both $A_\Phi(x_1) = a_1$ and $A_\Phi(x_2) = a_2$, thus by definition at least one of $\mathsf{f}_{B_1}(x_1, a_1)$ and $\mathsf{f}_{B_2}(x_2, a_2)$ is F, hence $\mathsf{f}_{B_1}, \mathsf{f}_{B_2}$ are consistent and F_1, F_2 as well. ■

4.2 Soundness

We next proceed to the the heart, and most technical part, of the proof of correctness, proving the construction is *sound*, that is, that in case Φ is far from being satisfiable, $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ has only a small independent set.

Lemma 4.4 (Soundness) *If $\alpha(\mathcal{G}_{[p,p^\bullet]}(\Phi)) \geq p^\bullet + \varepsilon$ then $\Upsilon_h(\Phi) \geq \varepsilon_\Phi$.*

Proof Overview: Assuming an independent set $\mathcal{I} \subset V$ of weight $\Lambda(\mathcal{I}) \geq p^\bullet + \varepsilon$, we consider for each block $B \in \mathcal{Z}$, its supposed long-code: the family $\mathcal{I}[B] = \mathcal{I} \cap V[B]$.

The first step (Lemma 4.6) is to find, for a non-negligible fraction of the blocks $\mathcal{Z}_q \subseteq \mathcal{Z}$, a small core of permissible block-assignments, and in it, one distinguished block-assignment to be used later to form an h-assignment satisfying ε_Φ of Φ. This is done by showing that for every $B \in \mathcal{Z}_q$, $\mathcal{I}[B]$ has both significant weight and low average-sensitivity. This, not necessarily true for p, is asserted for some slightly shifted value $q \in (p, p + \gamma)$. Utilizing Friedgut's lemma, we deduce the existence of a small core for $\mathcal{I}[B]$. Then, utilizing an Erdős-Ko-Rado-type bound on the maximal size of a 2-intersecting family, we find a distinguished block-assignment for each $B \in \mathcal{Z}_q$.

The next step is to focus on a subset of the blocks in $\mathcal{Z}_q$, that would serve as good representatives of $Z = X \times R_X$, from which emerges an h-assignment that semi-satisfies at least an ε_Φ fraction of $\varphi \in \Phi$. This is done by taking a random sub-block $\hat{B} \in Z^{(l-1)}$, and considering all blocks in $\mathcal{Z}$ that extend it. The distinguished block-assignment of each of these blocks form a multi-assignment $A_\mathcal{I}$ for Φ.

The final, most delicate part of the proof is Lemma 4.10, asserting that the distinguished block-assignments of the blocks extending $\hat{B}$ must be consistent. Indeed, since they all share the same $(l-1)$-sub-block $\hat{B}$, the consistency constraints these blocks impose on one another will be sufficiently tight. More accurately, we show there cannot be even h blocks extending $\hat{B}$ whose distinguished block-assignments are pairwise inconsistent, sufficing to conclude the proof.

Proof: Let then $\mathcal{I} \subset V$ be an independent set of size $\Lambda(\mathcal{I}) \geq p^\bullet + \varepsilon$, and denote, for each $B \in \mathcal{Z}$,

$$\mathcal{I}[B] \overset{def}{=} \mathcal{I} \cap V[B] .$$

The fractional size of $\mathcal{I}[B]$ within $V[B]$, according to Λ_B, is $\Lambda_B(\mathcal{I}[B]) = \mu_p(\mathcal{I}[B])$.

Assume w.l.o.g. that $\mathcal{I}$ is maximal, thus $\mathcal{I}[B]$, for any $B \in \mathcal{Z}$, is monotone and intersecting: It is intersecting, as $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ has edges connecting vertices corresponding to non-intersecting subsets, and it is monotone due to maximality:

Proposition 4.5 *Let $\mathcal{I}$ be an independent set of $\mathcal{G}_{[p,p^\bullet]}(\Phi)$. If $F \in \mathcal{I} \cap V[B]$, and $F \subset F' \in V[B]$, then $\mathcal{I} \cup \{F'\}$ is also an independent set.* ■

The first step in our proof is to find for a significant fraction of the blocks, a small core, and in it one distinguished block-assignment. Recall from Definition 2.9, that an element $\mathsf{f} \in C$ would be distinguished for a family $\langle \mathcal{I}[B] \rangle_C$ if there are two subsets $F^\flat, F^\sharp \in \langle \mathcal{I}[B] \rangle_C$ whose intersection is exactly $F^\flat \cap F^\sharp = \{\mathsf{f}\}$.

Friedgut's Lemma asserts the existence of a small core only for families with low average-sensitivity. We overcome this by slightly increasing p,

Lemma 4.6 *There exists some $q \in (p, p_{max})$, and a set of blocks $\mathcal{Z}_q \subseteq \mathcal{Z}$ whose size is $|\mathcal{Z}_q| \geq \frac{1}{4}\varepsilon \cdot |\mathcal{Z}|$, such that for all $B \in \mathcal{Z}_q$,*

1. *$\mathcal{I}[B]$ has an $(\frac{1}{16}\varepsilon, q)$-core, $\mathsf{Core}[B] \subset R_B$, of size $|\mathsf{Core}[B]| \leq h_0$.*
2. *The core-family $\langle \mathcal{I}[B] \rangle_{\mathsf{Core}[B]}$ has a distinguished element $\mathsf{f}^\diamondsuit[B] \in \mathsf{Core}[B]$.*

Proof: We will find a set of blocks $\mathcal{Z}_q \subseteq \mathcal{Z}$ such that for every $B \in \mathcal{Z}_q$, $\mathcal{I}[B]$ is of large weight and low average sensitivity, according to μ_q. We will then proceed to show that this implies the above properties. First consider blocks whose intersection with $\mathcal{I}$ has weight not much lower than the expectation,

$$\mathcal{Z}' \stackrel{def}{=} \left\{ B \in \mathcal{Z} \;\middle|\; \Lambda_B(\mathcal{I}[B]) > p^\bullet + \frac{1}{2}\varepsilon \right\}$$

By simple averaging, it follows that $|\mathcal{Z}'| \geq \frac{1}{2}\varepsilon \cdot |\mathcal{Z}|$, as otherwise

$$\Lambda(\mathcal{I}) \cdot |\mathcal{Z}| = \sum_{B \in \mathcal{Z}} \Lambda_B(\mathcal{I}[B]) \leq \frac{1}{2}\varepsilon\,|\mathcal{Z}| + \sum_{B \notin \mathcal{Z}'} \Lambda_B(\mathcal{I}[B]) < \frac{1}{2}\varepsilon\,|\mathcal{Z}| + \sum_{B \notin \mathcal{Z}'} (p^\bullet + \frac{1}{2}\varepsilon) \leq (p^\bullet + \varepsilon) \cdot |\mathcal{Z}|$$

Since μ_p is non-decreasing with p (see Proposition 2.4), and since the value of γ was chosen so that for every $q \in (p, p+\gamma)$, $p^\bullet + \frac{1}{4}\varepsilon > q^\bullet$, we have for every block $B \in \mathcal{Z}'$,

$$\mu_q(\mathcal{I}[B]) \geq \mu_p(\mathcal{I}[B]) > p^\bullet + \frac{1}{2}\varepsilon > q^\bullet + \frac{1}{4}\varepsilon\,. \qquad (*)$$

The family $\mathcal{I}[B]$, being monotone, cannot have high average sensitivity for many values of q, so by allowing an increase of at most γ, the set

$$\mathcal{Z}_q \stackrel{def}{=} \left\{ B \in \mathcal{Z}' \;\middle|\; \mathbf{as}_q(\mathcal{I}[B]) \leq \frac{2}{\gamma} \right\}$$

must be large for some $q \in (p, p+\gamma)$:

Proposition 4.7 *There exists $q \in (p, p+\gamma)$ so that $|\mathcal{Z}_q| \geq \frac{1}{4}\varepsilon \cdot |\mathcal{Z}|$.*

Proof: Consider the average, within $\mathcal{Z}'$, of the size of $\mathcal{I}[B]$ according to μ_q

$$\mu_q[\mathcal{Z}'] \stackrel{def}{=} |\mathcal{Z}'|^{-1} \cdot \sum_{B \in \mathcal{Z}'} \mu_q(\mathcal{I}[B])$$

and apply a version of Lagrange's Mean-Value Theorem. The derivative of $\mu_q[\mathcal{Z}']$ as a function of q is

$$\frac{d\mu_q[\mathcal{Z}']}{dq} = |\mathcal{Z}'|^{-1} \cdot \sum_{B \in \mathcal{Z}'} \frac{d\mu_q}{dq}(\mathcal{I}[B]) = |\mathcal{Z}'|^{-1} \cdot \sum_{B \in \mathcal{Z}'} \mathbf{as}_q(\mathcal{I}[B])$$

where the last equality follows from the Russo-Margulis identity (Lemma 2.5). Therefore, there must be some $q \in (p, p+\gamma)$ for which $\frac{d\mu_q[\mathcal{Z}']}{dq} \leq \frac{1}{\gamma}$, as otherwise $\mu_q[\mathcal{Z}']$ would increase too rapidly and $\mu_{p+\gamma}[\mathcal{Z}']$ would be larger than 1 which is impossible. It follows that at least half of the blocks in $\mathcal{Z}'$ have $\mathbf{as}_q(\mathcal{I}[B]) \leq \frac{2}{\gamma}$. ■

Fix then $q \in (p, p+\gamma)$, to be as in the proposition above, so that $|\mathcal{Z}_q| \geq \frac{1}{4}\varepsilon \cdot |\mathcal{Z}|$. We next show that the properties claimed by the lemma, indeed hold for all blocks in $\mathcal{Z}_q$.

The first property, namely that $\mathcal{I}[B]$ has an $(\frac{1}{16}\varepsilon, q)$-core, denoted $\mathsf{Core}[B] \subset R_B$, of size $|\mathsf{Core}[B]| \leq h_0$, is immediate from Friedgut's Lemma (see Theorem 2.3), plugging in the average sensitivity of $\mathcal{I}[B]$, and by definition of $h_0 = \sup_{q\in[p,p_{max}]} \mathsf{h}(q, \frac{1}{16}\,\varepsilon, \frac{2}{\gamma})$.

Denote the core-family approximating $\mathcal{I}[B]$ on $\mathsf{Core}[B]$, (see Definition 2.5), by $\mathcal{CF}_B \overset{def}{=} \langle\mathcal{I}[B]\rangle_{\mathsf{Core}[B]}$. By Proposition 2.10, since $\mathcal{I}[B]$ is monotone and intersecting, so is $\mathcal{CF}_B$. Moreover, Lemma 2.1 (a corollary of Friedgut's Lemma) asserts that

$$\mu_q(\mathcal{CF}_B) > \mu_q(\mathcal{I}[B]) - 3\cdot\frac{\varepsilon}{16} > q^\bullet$$

where the second inequality follows from inequality $(*)$ above. We can now utilize the bound on the maximal size of a 2-intersecting family (see Lemma 2.13), to deduce that $\mathcal{CF}_B$ is too large to be 2-intersecting, and must contain a distinguished element $\{\mathsf{f}^\diamondsuit\} \in \mathsf{Core}[B]$, and two subsets $F^\sharp, F^\flat \in \mathcal{CF}_B$ that intersect on exactly that block-assignment, $F^\sharp \cap F^\flat = \{\mathsf{f}^\diamondsuit\}$. This completes the proof of Lemma 4.6. ■

Let us now fix q as guaranteed by Lemma 4.6 above. The following implicit definitions appeared in the above proof, and will be used later as well,

Definition 4.2 (Core, Core-Family, Distinguished Block-Assignment) *Let $B \in \mathcal{Z}_q$.*

- *B's* core, *denoted $\mathsf{Core}[B] \subset R_B$, is an arbitrary smallest $(\frac{1}{16}\varepsilon, q)$-core of $\mathcal{I}[B]$.*
- *B's* core-family, *is the core-family on B's core (see Definition 2.5), denoted $\mathcal{CF}_B = \langle\mathcal{I}[B]\rangle_{\mathsf{Core}[B]}$.*
- *B's* distinguished block-assignment, *is an arbitrary distinguished element of $\mathcal{CF}_B$, denoted $\mathsf{f}^\diamondsuit[B] \in \mathsf{Core}[B]$; i.e. for which there exist $F^\sharp, F^\flat \in \mathcal{CF}_B$ with $F^\sharp \cap F^\flat = \{\mathsf{f}^\diamondsuit[B]\}$.*

Let us further define for each block $B \in \mathcal{Z}_q$, the set of all block-assignments of B that have non-negligible influence:

Definition 4.3 (Extended Core) *For $B \in \mathcal{Z}$, let the* extended core *of B be*

$$\mathit{ECore}[B] \overset{def}{=} \mathit{Core}[B] \cup \left\{\mathsf{f} \in R_B \;\middle|\; \mathbf{influence}^{\mathsf{f}}_q(\mathcal{I}[B]) \geq \eta\right\}$$

The average-sensitivity of $\mathcal{I}[B]$ is defined to be the sum of the elements' influences, thus, since for every $B \in \mathcal{Z}_q$, it is bounded by $\mathbf{as}_q(\mathcal{I}[B]) \leq \frac{2}{\gamma}$,

$$|\mathsf{ECore}[B]| \leq \frac{\mathbf{as}_q(\mathcal{I}[B])}{\eta} + h_0 \leq \left\lceil \frac{2}{\gamma\cdot\eta}\right\rceil + h_0 = h_1$$

The next step in our proof, is to identify an $(l-1)$-sub-block $\hat{B} \in Z^{(l-1)}$ whose extensions $\hat{B} \cup \{z\}$ represent $Z = X \times R_X$, and whose distinguished block-assignments will be sufficiently consistent for defining an h-assignment for Φ. Analyzing the consistency between the distinguished block-assignments of distinct blocks, is complicated by the fact that families encoding distinct blocks consist of subsets of distinct domains ($R_{B_1} \neq R_{B_2}$ for $B_1 \neq B_2$). Considering only the blocks that extend a specific sub-block $\hat{B} \in Z^{(l-1)}$, yields a nice 2-to-2 correspondence between their block-assignments. The block-assignments of blocks $B = \hat{B} \cup \{z\}$ are paired according to their restriction to $\hat{B}$, such that all the pairs whose restriction is mapped to the same sub-block-assignment naturally correspond to each other.

It would be undesired to have both block-assignments in a given pair influential in $\mathcal{I}[B]$ for this would mean that the structure of $\mathcal{I}[B]$ is not preserved when reduced to $\hat{B}$. Thus, besides requiring that many of the blocks $\hat{B} \cup \{z\}$ extending $\hat{B}$ reside in $\mathcal{Z}_q$, we need them to be *preserved* by $\hat{B}$:

Definition 4.4 (Preservation) *Let $B \in \mathcal{Z}$, and let $\hat{B} \subset B$, $|\hat{B}| = l-1$. Let us denote by $\mathsf{f}|_{\hat{B}}$ the restriction to $\hat{B}$ of a block-assignment $\mathsf{f} \in R_B$. We say that $\hat{B}$* preserves *B, if there is no pair of block-assignments $\mathsf{f}_1 \neq \mathsf{f}_2 \in R_B$ with $\mathsf{f}_1|_{\hat{B}} = \mathsf{f}_2|_{\hat{B}}$, such that $\mathsf{f}_1, \mathsf{f}_2 \in \mathit{ECore}[B]$.*

It is almost always the case that $\hat{B}$ preserves $\hat{B} \cup \{z\}$:

Proposition 4.8

$$\forall B \in \mathcal{Z} \quad |\{\, z \in B \mid B \setminus \{z\} \textit{ does not preserve } B\}| < \frac{(h_1)^2}{2}\,.$$

Proof: Each pair of block-assignments $\mathsf{f}_1, \mathsf{f}_2 \in \mathsf{ECore}[B]$ can cause at most one $\hat{B}$ to not preserve B, and for any block $B \in \mathcal{Z}_q$, $|\mathsf{ECore}[B]| \leq h_1$; consequently, the number of $\hat{B}$ not preserving B is at most $\binom{h_1}{2} < \frac{(h_1)^2}{2}$. ■

The last step before identifying the required $\hat{B}$ is to note that a distinguished block-assignment for a block $\hat{B} \cup \{z\}$ is useful for constructing an assignment for Φ, if it assigns T to $z = (x, a)$. Hence, for each $\hat{B}$ we consider the following set $Z_{\hat{B}} \subset Z$:

Definition 4.5 *Let $Z_{\hat{B}} \subseteq Z$ be:*

$$Z_{\hat{B}} \overset{def}{=} \left\{ z \in Z \setminus \hat{B} \;\middle|\; B = \hat{B} \cup \{z\} \in \mathcal{Z}_q, \quad \textit{and} \quad \hat{B} \textit{ preserves } B, \quad \textit{and} \quad \mathsf{f}^{\diamondsuit}[B](z) = \mathsf{T} \right\}$$

It follows from the definition of $Z_{\hat{B}}$, that if $z_1, z_2 \in Z_{\hat{B}}$ are Z-inconsistent (recall Definition 3.3), then the distinguished block-assignments of $B_1 = \hat{B} \cup \{z_1\}$ and $B_2 = \hat{B} \cup \{z_2\}$ are inconsistent, $\mathsf{f}^{\diamondsuit}[B_1] \bowtie \mathsf{f}^{\diamondsuit}[B_2]$, (see Definition 3.5). Finally, let us identify a sub-block $\hat{B}$, for which $Z_{\hat{B}}$ is large:

Proposition 4.9 *There exists $\hat{B} \in \binom{Z}{l-1}$, with $|Z_{\hat{B}}| \geq \frac{1}{32}\varepsilon\frac{|Z|}{|R_X|} = \frac{1}{32}\varepsilon\,|X|$.*

Proof: We write

$$\Pr_{\hat{B},\, z\in Z\setminus\hat{B}}\left[z \in Z_{\hat{B}}\right] \geq \frac{1}{4}\varepsilon \cdot \Pr_{B,\, z\in B}\left[z \in Z_{B\setminus\{z\}} \,\middle|\, B \in \mathcal{Z}_q\right] \geq \frac{1}{4}\varepsilon \cdot \frac{1}{4\,|R_X|}$$

where the first inequality follows from Proposition 4.7 asserting $\mathcal{Z}_q \geq \frac{1}{4}\varepsilon\,|\mathcal{Z}|$. The second inequality is a consequence of the fact that for any $\mathsf{f} \in R_B$, there are at least $t = \frac{l}{2|R_X|}$ elements $z \in B$ with $\mathsf{f}(z) = \mathsf{T}$; and at most $\frac{(h_1)^2}{2}$ $(l-1)$-blocks $\hat{B} \subset B$ not preserving B; hence, conditioned on $B \in \mathcal{Z}_q$, the probability of $z \in Z_{\hat{B}}$ is at least $\frac{1}{2|R_X|} - \frac{(h_1)^2}{2l} \geq \frac{1}{4|R_X|}$ as $l \geq 2(h_1)^2 \cdot |R_X|$.

There must therefore be at least one $\hat{B}$ for which $\Pr_{z\in Z\setminus\hat{B}}\left[z \in Z_{\hat{B}}\right] \geq \frac{\varepsilon}{16|R_X|}$, hence, $|Z_{\hat{B}}| \geq \frac{1}{16|R_X|}\varepsilon \cdot \left|Z \setminus \hat{B}\right| \geq \frac{1}{32}\varepsilon \cdot |X|$, as $l < \frac{1}{2}\,|X|$. ■

The h-Assignment $A_{\mathcal{I}}$. We are now ready to present, based on $Z_{\hat{B}}$, an h-assignment $A_{\mathcal{I}}$ semi-satisfying Φ. Fix one such $\hat{B}$ for which $|Z_{\hat{B}}| \geq \frac{1}{32}\varepsilon \cdot |X|$. For every $x \in X$ and $y \in Y$, set

$$\begin{aligned} A_{\mathcal{I}}(x) &\overset{def}{=} \{a \in R_X \mid (x,a) \in Z_{\hat{B}}\} \\ A_{\mathcal{I}}(y) &\overset{def}{=} \bigcup_{\varphi_{x\to y}\in\Phi} \varphi_{x\to y}(A_{\mathcal{I}}(x)) \end{aligned}$$

Ideally, we would have liked to rule out the possibility of two z_i's in $Z_{\hat{B}}$ being Z-inconsistent, in which case for every x, y, $|A_{\mathcal{I}}(x)|, |A_{\mathcal{I}}(y)| \leq 1$. Instead, we show there cannot be h values in $Z_{\hat{B}}$ that are pairwise Z-inconsistent. Specifically, we prove,

Lemma 4.10 *For every $x \in X$, $|A_{\mathcal{I}}(x)| < h$, and for every $y \in Y$, $|A_{\mathcal{I}}(y)| < h$.*

This lemma implies that $\Upsilon_h(\Phi) \geq \varepsilon_\Phi$, due to the following. By the definition of $A_{\mathcal{I}}$, for every x with $A_{\mathcal{I}}(x) \neq \phi$ and for *every* $\varphi_{x\to y} \in \Phi$,

$$\varphi_{x\to y}(A_{\mathcal{I}}(x)) \cap A_{\mathcal{I}}(y) \neq \phi\,.$$

Denote $X_{\hat{B}} = \{x \in X \mid A_{\mathcal{I}}(x) \neq \phi\}$ and observe that, since Φ is regular (see section 3.1), there is an equal number of $\varphi_{x\to y} \in \Phi$ for each variable x, therefore $A_{\mathcal{I}}$'s success probability is:

$$\Pr_{\varphi_{x\to y}\in\Phi}\left[\varphi_{x\to y}(A_{\mathcal{I}}(x)) \cap A_{\mathcal{I}}(y) \neq \phi\right] = \frac{|X_{\hat{B}}|}{|X|} > \frac{1}{h}\frac{|Z_{\hat{B}}|}{|X|} \geq \frac{\varepsilon}{32h}\frac{|X|}{|X|} = \varepsilon_\Phi\,.$$

Proof: (of Lemma 4.10) Assume, by way of contradiction, that there exist $z_1, \ldots, z_h \in Z_{\hat{B}}$, where $B_i = \hat{B} \cup \{z_i\}$, and such that z_i are pairwise Z-inconsistent; we would then show that $\cup_{i\in[h]}\mathcal{I}[B_i]$ is not an independent set.

This suffices so as to prove the lemma, as, a variable $x \in X$, with $|A_{\mathcal{I}}(x)| \geq h$, or $y \in Y$, with $|A_{\mathcal{I}}(y)| \geq h$, implies h blocks $B_i = \hat{B} \cup \{z_i\}$ with pairwise Z-inconsistent z_i's.

Let us begin with a brief sketch of the proof that follows. Analyzing consistency between blocks $\hat{B} \cup \{z_i\}$ leads us to consider the common sub-block $\hat{B}$, and the sub-block-assignments that are restrictions of block-assignments in R_{B_i} to $\hat{B}$.

We begin by defining these restrictions, for the core, the distinguished block assignment, and the extended core of each block. Next we find – applying some combinatorics (namely, a pigeon-hole principle and a sunflower lemma) – a pair of blocks B_1 and B_2, out of the h, whose encodings $\mathcal{I}[B_1], \mathcal{I}[B_2]$ are 'extremely-inconsistent'. We then proceed in a series of steps (Propositions 4.12–4.16) to identify a pair of subsets $F_1 \in \mathcal{I}[B_1]$ and $F_2 \in \mathcal{I}[B_2]$ with an edge between them.

The $(l-1)$-*block-assignments* of $\hat{B} \in \binom{Z}{l-1}$, are defined to be

$$R_{\hat{B}} \overset{def}{=} \left\{ \mathsf{f} \colon \hat{B} \to \{\mathsf{T}, \mathsf{F}\} \right\}$$

A block-assignment $\mathsf{f} \in R_{B_i}$ has a natural restriction to $\hat{B}$, denoted $\mathsf{f}|_{\hat{B}} \in R_{\hat{B}}$, where $\mathsf{f}|_{\hat{B}}(z) = \mathsf{f}(z)$.

For the remaining analysis, let us name the three important entities regarding each block B_i, for $i \in [h]$: B_i's distinguished block-assignment, the core of B_i, and the extended core of B_i,

$$\mathsf{f}_i^{\diamondsuit} \overset{def}{=} \mathsf{f}^{\diamondsuit}[B_i] \qquad C_i \overset{def}{=} \mathsf{Core}[B_i] \qquad E_i \overset{def}{=} \mathsf{ECore}[B_i]$$

and their natural restrictions to $\hat{B}$ (where the natural restriction of a set is the set comprising the retrictions of its elements),

$$\hat{\mathsf{f}}_i^{\diamondsuit} \overset{def}{=} \mathsf{f}_i^{\diamondsuit}|_{\hat{B}} \qquad \hat{C}_i \overset{def}{=} C_i|_{\hat{B}} \qquad \hat{E}_i \overset{def}{=} E_i|_{\hat{B}}$$

Now, recall the core-family $\mathcal{CF}_{B_i}$, which is the family of subsets, over the core of each B_i, each of which extension is of $\frac{3}{4}$ weight in $\mathcal{I}[B_i]$. For each block B_i, $i \in [h]$, $\mathsf{f}_i^{\diamondsuit}$ being distinguished implies a pair of subsets

$$F_i^{\flat}, F_i^{\sharp} \in \mathcal{CF}_{B_i} \text{ so that } F_i^{\flat} \cap F_i^{\sharp} = \left\{ \mathsf{f}_i^{\diamondsuit} \right\}$$

Let their natural restriction to $\hat{B}$ be

$$\hat{F^{\flat}}_i \overset{def}{=} F_i^{\flat}|_{\hat{B}} \qquad \hat{F^{\sharp}}_i \overset{def}{=} F_i^{\sharp}|_{\hat{B}}$$

and note that, as $\hat{B}$ preserves every B_i, it follows that, for all $i \in [h]$,

$$\hat{F^{\flat}}_i \cap \hat{F^{\sharp}}_i = \left\{ \hat{\mathsf{f}}_i^{\diamondsuit} \right\} \tag{4.1}$$

Let us add some brief intuition for the first step of the proof. Our first goal is to identify two blocks B_{i_1} and B_{i_2} that are extremely inconsistent. This task would be easy had we found two blocks whose restricted extended-cores $\hat{E}_i$ are disjoint. Loosely speaking, this follows since anything that is of any remote influence in the encoding of a block, takes place within the extended-core. Resorting to the next best thing, to our aid comes a combinatorial lemma, that identifies a subset of the blocks, whose $\hat{E}_i$'s are pairwise disjoint, *except* for a common center. This is achieved by the Erdős-Rado Sunflower lemma:

Lemma 4.11 ([ER60]) *For any $\mathcal{F} \subset \binom{R}{k}$, if $|\mathcal{F}| \geq d^k \cdot k!$, there are d distinct sets $F_1, \ldots, F_d \in \mathcal{F}$, such that, let $\mathbf{\Delta} \overset{def}{=} F_1 \cap \ldots \cap F_d$, the sets $F_i \setminus \mathbf{\Delta}$ are pairwise disjoint.*

The sets $F_1, .., F_d$ are called a Sunflower, or a $\mathbf{\Delta}$-system. This statement can easily be extended to families in which each subset is of size *at most* k.

We apply this lemma for $R = R_{\hat{B}}$, and $\mathcal{F} = \{\hat{E}_1, .., \hat{E}_h\}$. Recall (definition 4.1) we have fixed $h = (h_s)^{h_1} \cdot (h_1)!$, hence Lemma 4.11 implies there exists some $J \subseteq [h]$, $|J| = h_s$, such that

$$\left\{\hat{E}_i \setminus \mathbf{\Delta}\right\}_{i \in J} \text{ are pairwise disjoint for } \mathbf{\Delta} \overset{def}{=} \bigcap_{i \in J} \hat{E}_i \tag{4.2}$$

Out of these h_s blocks, we will find, applying a pigeon-hole principle, a pair of inconsistent blocks. Inconsistent, in this context, means that the core-families of these blocks contain two subsets, $F_1 \in \mathcal{CF}_{B_1}$ and $F_2 \in \mathcal{CF}_{B_2}$, whose block-assignments are pairwise inconsistent. As the blocks B_i for $i \in J$ have pairwise disjoint $\hat{E}_i$s outside $\mathbf{\Delta}$, we need to consider only block-assignments whose restrictions fall into $\mathbf{\Delta}$.

Consider, for each $i \in J$, the triplet $\left\langle \hat{C}_i \cap \mathbf{\Delta}, \quad \hat{F^\flat}_i \cap \mathbf{\Delta}, \quad \hat{F^\sharp}_i \cap \mathbf{\Delta} \right\rangle$, and note that, since $\hat{F^\flat}_i, \hat{F^\sharp}_i \subseteq \hat{C}_i$ the number of possible triplets is at most

$$\begin{aligned}\left|\left\{ \left\langle \hat{C} \cap \mathbf{\Delta}, \hat{F^\flat} \cap \mathbf{\Delta}, \hat{F^\sharp} \cap \mathbf{\Delta} \right\rangle \;\middle|\; |\hat{C}| \leq h_0, \hat{F^\flat}, \hat{F^\sharp} \subseteq \hat{C} \right\}\right| \;&\leq\; \sum_{k=0}^{h_0} \binom{h_1}{k} \cdot 2^{h_0} \cdot 2^{h_0} \\ &<\; h_s = |J|\end{aligned}$$

(recall we have set (definition 4.1) $h_s = 1 + 2^{2h_0} \cdot \sum_{k=0}^{h_0} \binom{h_1}{k}$). Therefore, by the pigeon-hole principle, there must be some $i_1, i_2 \in J$ for which

$$\left\langle \hat{C}_{i_1} \cap \mathbf{\Delta},\ \hat{F^\flat}_{i_1} \cap \mathbf{\Delta},\ \hat{F^\sharp}_{i_1} \cap \mathbf{\Delta} \right\rangle = \left\langle \hat{C}_{i_2} \cap \mathbf{\Delta},\ \hat{F^\flat}_{i_2} \cap \mathbf{\Delta},\ \hat{F^\sharp}_{i_2} \cap \mathbf{\Delta} \right\rangle \tag{4.3}$$

Assume w.l.o.g. that $i_1 = 1$, $i_2 = 2$. We will arrive at a contradiction by finding an edge between the blocks B_1, B_2, specifically, by finding two extensions, one of $F_1^\flat$ in $\mathcal{I}[B_1]$, and another of $F_2^\sharp$ in $\mathcal{I}[B_2]$, all of whose block-assignments are pairwise inconsistent.

As a first step, let us prove that the block-assignments in $F_1^\flat$ and $F_2^\sharp$ are pairwise inconsistent:

Proposition 4.12

$$\mathsf{f}_1 \in F_1^\flat, \mathsf{f}_2 \in F_2^\sharp \quad \Rightarrow \mathsf{f}_1 \bowtie \mathsf{f}_2$$

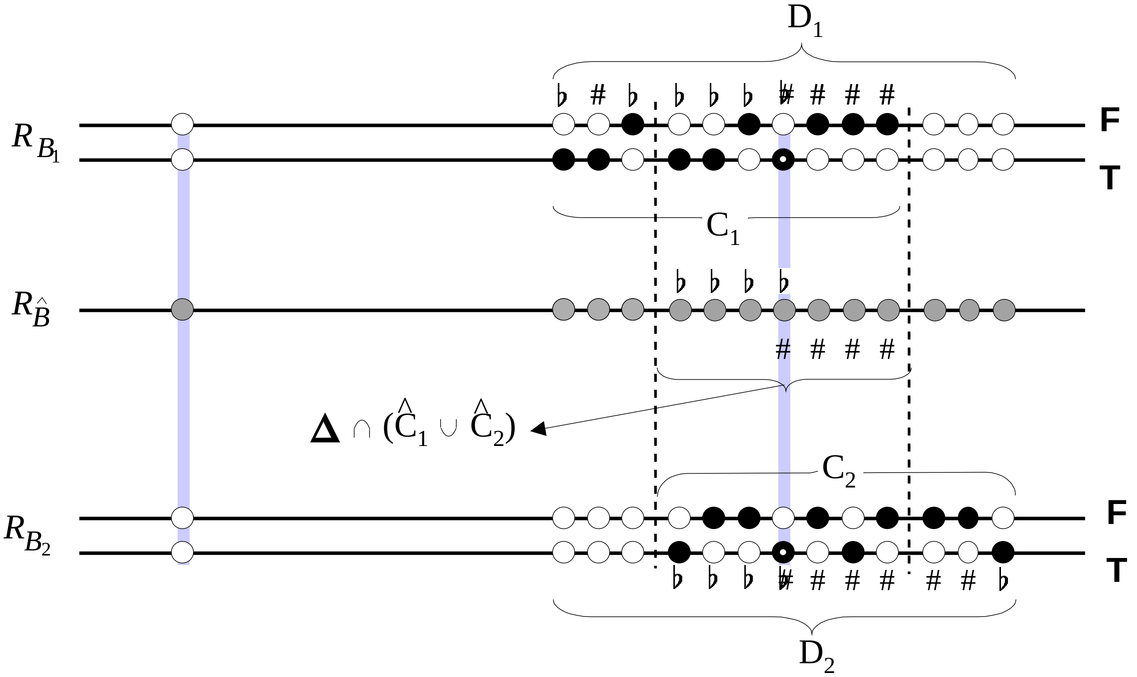

Figure 4.1: Block Assignments of B_1, B_2 and sub-block-assignments of $\hat{B}$. R_{B_1} (resp. R_{B_2}) is represented by the two upper (resp. two lower) horizontal lines labelled by T and F to indicate the value assigned to v_1 (resp. v_2) by block-assignments on that line. Each circle represents a single block assignment. On the left a column (highlighted as a light gray vertical line) consists of four block assignments and a sub-block assignment which is their common restriction to $\hat{B}$. All block assignments in the same column agree on their restriction to $\hat{B}$, depicted as a gray circle on the middle horizontal line that represents $R_{\hat{B}}$. Two block assignments are consistent *only* if they are in the same column and are not both T. The blackened circles represent members of the core of B_1 and the block-assignments in $F_1^\flat$ and $F_1^\sharp$ are labelled $\flat$ and $\sharp$. The distinguished block-assignment – marked by a white dot – is labelled by both $\flat$ and $\sharp$, and assigns T to v_1. The dashed vertical lines border the intersection of $\hat{C}_1$ with $\hat{C}_2$, which is equal to $\hat{C}_1 \cap \hat{C}_2 = \hat{C}_1 \cap \mathbf{\Delta} = \hat{C}_2 \cap \mathbf{\Delta}$ and is where the restrictions of $F_1^\sharp, F_1^\flat$ are equal to those of $F_2^\sharp, F_2^\flat$. This also implies (see Proposition 4.15) that $(\hat{D}_1 \setminus \hat{C}_1) \cap \mathbf{\Delta} \subseteq (\hat{D}_1 \setminus \hat{C}_1) \cap \hat{E}_1 = \phi$.

Proof: For f_1 to be consistent with f_2, it must be that $\mathsf{f}_1|_{\hat{B}} = \mathsf{f}_2|_{\hat{B}} \in \hat{F_1^\flat} \cap \hat{F_2^\sharp} \subseteq \hat{E}_1 \cap \hat{E}_2 = \mathbf{\Delta}$. B_1 and B_2 are chosen (see equation 4.3) so that $\hat{F_1^\flat} \cap \mathbf{\Delta} = \hat{F^\flat}_2 \cap \mathbf{\Delta}$ and $\hat{F_2^\sharp} \cap \mathbf{\Delta} = \hat{F^\sharp}_1 \cap \mathbf{\Delta}$. Consequently $\mathsf{f}_1|_{\hat{B}} = \mathsf{f}_2|_{\hat{B}} \in \hat{F_1^\flat} \cap \hat{F^\sharp}_1 \cap \mathbf{\Delta} = \hat{F^\flat}_2 \cap \hat{F_2^\sharp} \cap \mathbf{\Delta}$, however, equation (4.1) asserts that the only block-assignment in these two intersections is the distinguished one, hence $\hat{\mathsf{f}}_1^{\Diamond} = \mathsf{f}_1|_{\hat{B}} = \mathsf{f}_2|_{\hat{B}} = \hat{\mathsf{f}}_2^{\Diamond}$. Since $\hat{B}$ preserves both B_1 and B_2, $\mathsf{f}_1 = \mathsf{f}_1^{\Diamond}$ and $\mathsf{f}_2 = \mathsf{f}_2^{\Diamond}$. However, $\mathsf{f}_1^{\Diamond} \bowtie \mathsf{f}_2^{\Diamond}$, as they assign T to both z_1 and z_2 that are Z-inconsistent. ■

It may well be that $F_1^\flat \notin \mathcal{I}[B_1]$ and $F_2^\sharp \notin \mathcal{I}[B_2]$, thus the fact that they are inconsistent (and so, connected by an edge) is only a first step towards a contradiction. Nevertheless, we know that $F_1^\flat \in \mathcal{CF}_{B_1} = \langle \mathcal{I}[B_1] \rangle_{\mathsf{Core}[B_1]}$ means that $\frac{3}{4}$ of $\left\{F_1^\flat\right\} \sqcup \mathbf{P}\left(R_{B_1} \setminus \mathsf{Core}[B_1]\right)$ are in $\mathcal{I}[B_1]$; and likewise for $F_2^\sharp$. In what follows, we utilize this large volume of $\frac{3}{4}$ to find extensions of these sets, that are in $\mathcal{I}$, yet are inconsistent.

Let us partition the set of $(l-1)$-block assignments of $R_{\hat{B}}$ into the important ones, which are restrictions of block-assignments in the cores of B_1 or B_2, and the rest,

$$\hat{D} = \hat{C}_1 \cup \hat{C}_2 \quad \text{and} \quad \hat{R} = R_{\hat{B}} \setminus \hat{D}$$

which immediately partitions the block-assignments of R_{B_1} and R_{B_2}, according to whether their restriction falls within $\hat{D}$:

$$D_1 = \left\{ \mathsf{f} \in R_{B_1} \;\middle|\; \mathsf{f}|_{\hat{B}} \in \hat{D} \right\} \quad \text{and} \quad R_1 = R_{B_1} \setminus D_1$$

and similarly for R_{B_2},

$$D_2 = \left\{ \mathsf{f} \in R_{B_2} \;\middle|\; \mathsf{f}|_{\hat{B}} \in \hat{D} \right\} \quad \text{and} \quad R_2 = R_{B_2} \setminus D_2$$

Proposition 4.13 $|D_1| \le 4h_0$ *and* $|D_2| \le 4h_0$.

Proof: Simply note that $|D_1|, |D_2| \le 2|\hat{D}| \le 2(|\hat{C}_1| + |\hat{C}_2|) \le 2(|C_1| + |C_2|) = 4h_0$. ■

So far we have established two subsets $F_1^\flat \in \mathcal{CF}_{B_1}$ and $F_2^\sharp \in \mathcal{CF}_{B_2}$ in the core-families of B_1 and B_2, all of whose block-assignments are pairwise inconsistent. Furthermore, $F_1^\flat \in \mathbf{P}(C_1) \subseteq \mathbf{P}(D_1)$ and $F_2^\sharp \in \mathbf{P}(C_2) \subseteq \mathbf{P}(D_2)$, hence it suffices to show two subsets $H_1 \in \mathbf{P}(R_1)$ and $H_2 \in \mathbf{P}(R_2)$ all of whose block-assignments are pairwise-inconsistent, and so that $F_1^\flat \cup H_1 \in \mathcal{I}[B_1]$ and $F_2^\sharp \cup H_2 \in \mathcal{I}[B_2]$.

Let us prove this by showing first that the families of subsets extending $F_1^\flat$ and $F_2^\sharp$ within $\mathcal{I}$ are large; and then proceed to show that there are two subsets, H_1 and H_2 as required.

Let us first name these two families of subsets extending $F_1^\flat$ and $F_2^\sharp$ within $\mathcal{I}$:

$$\mathcal{I}_1 = \left\{ F \in \mathbf{P}(R_1) \;\middle|\; (F_1^\flat \cup F) \in \mathcal{I}[B_1] \right\} \quad \text{and} \quad \mathcal{I}_2 = \left\{ F \in \mathbf{P}(R_2) \;\middle|\; (F_2^\sharp \cup F) \in \mathcal{I}[B_2] \right\}$$

and proceed to prove they are large:

Proposition 4.14

$$\mu_q^{R_1}(\mathcal{I}_1) > \frac{1}{2} \qquad \textit{and} \qquad \mu_q^{R_2}(\mathcal{I}_2) > \frac{1}{2}$$

Proof: Let us prove the first case; the second one is proven by a symmetric, but otherwise identical, argument. By definition of $\mathcal{CF}_{B_1} = \langle \mathcal{I}[B_1] \rangle_{C_1}$, it is the case that

$$\Pr_{F \in \mu_q} \left[F \in \mathcal{I}[B_1] \;\; \middle| \; F \cap C_1 = F_1^\flat \right] > \frac{3}{4}$$

Note that the only difference between this event and

$$\mu_q^{R_1}(\mathcal{I}_1) = \Pr_{F \in \mu_q} \left[F \in \mathcal{I}[B_1] \;\; \middle| \; F \cap D_1 = F_1^\flat \right]$$

is the conditioning on F to not contain any block-assignment in $D_1 \setminus C_1$. Simplistically, if the elements in $D_1 \setminus C_1$ have tiny influence, then removing them from a subset does not take it out of $\mathcal{I}[B_1]$. Hence, it suffices to prove that this family, of extensions of $F_1^\flat$ within $\mathcal{I}[B_1]$, is almost independent of the set of block-assignments $D_1 \setminus C_1$, that is, that one can extract a small ($< \frac{1}{4}$) fraction of $\mathcal{I}_1$ and make it completely independent of the block-assignments outside $R_1 \cup C_1$.

Let us first observe that block-assignments in $D_1 \setminus C_1$ indeed have tiny influence,

Proposition 4.15

$$(D_1 \setminus C_1) \cap E_1 = \phi$$

Proof: There are two cases to consider for $\mathsf{f} \in D_1 \setminus C_1$: Either $\mathsf{f}|_{\hat{B}} \in \hat{C}_1$ and in that case, since $\hat{B}$ preserves B_1 and since $\mathsf{f} \notin C_1$, $\mathsf{f} \notin E_1$; or, $\mathsf{f}|_{\hat{B}} \in \hat{C}_2 \setminus \hat{C}_1$ and since B_1 and B_2 are chosen (¶) so that $\hat{C}_1 \cap \mathbf{\Delta} = \hat{C}_2 \cap \mathbf{\Delta}$, we deduce $\mathsf{f}|_{\hat{B}} \notin \mathbf{\Delta}$. Now $\mathsf{f}|_{\hat{B}} \in \hat{C}_2 \subseteq \hat{E}_2$, implies $\mathsf{f}|_{\hat{B}} \notin \hat{E}_1$, thus $\mathsf{f} \notin E_1$. ∎

By definition of the extended core E_i (Definition 4.2), it follows that for every $\mathsf{f} \in D_1 \setminus C_1$, $\textbf{influence}_q^{\mathsf{f}}(\mathcal{I}[B_1]) < \eta$. Since $|D_1 \setminus C_1| < 4h_0$ (Proposition 4.13) we can deduce that $\mathcal{I}[B_1]$ is almost independent of $D_1 \setminus C_1$, utilizing a relatively simple, general property related to influences. Namely, that, given any family of subsets of a domain R, and a set $T \subset R$ of elements of tiny influence, one has to remove only a small fraction of the family to make it completely independent of T. More accurately, Proposition 2.8 asserts that if the influence of elements in $D_1 \setminus C_1$ is bounded by η, then the weight of the subsets in $\mathcal{I}[B_1]'$, i.e. those that have to be removed from $\mathcal{I}[B_1]$ to make it independent of $D_1 \setminus C_1$,

$$\mathcal{I}[B_1]' \overset{def}{=} \left\{ F \in \mathcal{I}[B_1] \;\middle|\; (F \setminus (D_1 \setminus C_1)) \notin \mathcal{I}[B_1] \right\},$$

is bounded by

$$\mu_q^{R_{B_1}}(\mathcal{I}[B_1]') < 4h_0 \cdot \eta \cdot q^{-4h_0} \le \frac{1}{4} q^{4h_0}$$

since $\eta = \frac{1}{16h_0} \cdot p^{8h_0} \leq \frac{1}{16h_0} \cdot q^{8h_0}$, see Definition 4.1.

Even if all $\mathcal{I}[B_1]'$ is concentrated on $F_1^\flat$, since $F_1^\flat$'s weight in $\mathbf{P}(D_1)$ is at least $q^{|D_1|} \geq q^{4h_0}$, $\mu_q^{D_1}\left(F_1^\flat\right) \geq q^{4h_0}$, it follows that (using $\Pr(A \mid B) \leq \Pr(A)/\Pr(B)$),

$$\Pr_{F \in \mu_q^{R_1}}\left[F \in \mathcal{I}[B_1]' \mid F \cap D_1 = F_1^\flat\right] \leq \Pr_{F \in \mu_q^{R_1}}\left[F \in \mathcal{I}[B_1]'\right] \cdot \frac{1}{\mu_q^{D_1}\left(F_1^\flat\right)} < \frac{1}{4}$$

Proposition 4.14 is thereby proven. ■

We complete the proof of the Soundness Lemma, by deducing from the large volume of I_1, I_2, the existence of two subsets $H_1 \in I_1$ and $H_2 \in I_2$ so that $\langle H_1, H_2 \rangle \in E$, implying $\left\langle F_1^\flat \cup H_1, F_2^\sharp \cup H_2 \right\rangle \in E$, which is the desired contradiction.

Proposition 4.16 *Let $I_1 \subset \mathbf{P}(R_1), I_2 \subset \mathbf{P}(R_2)$. If $(1-q)^2 \geq q$ and $\mu_q^{R_1}(\mathcal{I}_1) + \mu_q^{R_2}(\mathcal{I}_2) > 1$, there exist $H_1 \in \mathcal{I}_1$ and $H_2 \in \mathcal{I}_2$ such that $\langle H_1, H_2 \rangle \in E$.*

Proof: This proposition is proven by modifying the proof for the case of cross-intersecting families (Proposition 2.9). In that proof, we bounded the size of a pair of cross-intersecting families by pairing each subset with its complement, noting that at $p = \frac{1}{2}$ their weights are equal.

In this case, we focus on the value $q = p_{max} = \frac{3-\sqrt{5}}{2}$ for which $(1-q)^2 = q$, noting that since $q \leq p_{max}$, the monotonicity of I_1, I_2 (see Proposition 2.4) yields $\mu_{p_{max}}(I_1) + \mu_{p_{max}}(I_2) > 1$. Here let us partition both $\mathbf{P}(R_1)$ and $\mathbf{P}(R_2)$, and define an appropriate 'complement' for each part, rather than for each subset.

Our partition is defined according to a 'representative mapping' mapping each $F \in \mathbf{P}(R_1)$ to a function $\Pi[F_1] : \hat{R} \to \left\{\overline{\mathsf{TF}}, \mathsf{T}\overline{\mathsf{F}}, \mathsf{F}\right\}$ defined as follows:

$$\forall \hat{\mathsf{f}} \in \hat{R}, \quad \Pi[F_1](\hat{\mathsf{f}}) \overset{def}{=} \begin{cases} \overline{\mathsf{TF}} & \hat{\mathsf{f}}^{(z_1 \leftarrow \mathsf{T})}, \hat{\mathsf{f}}^{(z_1 \leftarrow \mathsf{F})} \notin F_1 \\ \mathsf{T}\overline{\mathsf{F}} & \hat{\mathsf{f}}^{(z_1 \leftarrow \mathsf{T})} \in F_1, \ \hat{\mathsf{f}}^{(z_1 \leftarrow \mathsf{F})} \notin F_1 \\ \mathsf{F} & \hat{\mathsf{f}}^{(z_1 \leftarrow \mathsf{F})} \in F_1 \end{cases}$$

(symmetrically, we define $\Pi[F_2]$ for each $F_2 \in \mathbf{P}(R_2)$). This mapping is natural when considering the characteristic function of F_1 and asking, for every $\hat{\mathsf{f}} \in \hat{R}$, the value of that function on the two extensions of $\hat{\mathsf{f}}$ in R_1, $\hat{\mathsf{f}}^{(z_1 \leftarrow \mathsf{T})}$ and $\hat{\mathsf{f}}^{(z_1 \leftarrow \mathsf{F})}$.

Additionally, for a function $\Pi = \Pi[F_1]$, $\Pi : \hat{R} \to \left\{\overline{\mathsf{TF}}, \mathsf{T}\overline{\mathsf{F}}, \mathsf{F}\right\}$, let its complement be $\Pi^{\mathsf{c}} : \hat{R} \to \left\{\overline{\mathsf{TF}}, \mathsf{T}\overline{\mathsf{F}}, \mathsf{F}\right\}$ defined as follows:

$$\forall \hat{\mathsf{f}} \in \hat{R}, \quad \Pi^{\mathsf{c}}(\hat{\mathsf{f}}) \overset{def}{=} \begin{cases} \overline{\mathsf{TF}} & \Pi(\hat{\mathsf{f}}) = \mathsf{F} \\ \mathsf{T}\overline{\mathsf{F}} & \Pi(\hat{\mathsf{f}}) = \mathsf{T}\overline{\mathsf{F}} \\ \mathsf{F} & \Pi(\hat{\mathsf{f}}) = \overline{\mathsf{TF}} \end{cases}$$

Observe that $\Pi^{\mathsf{cc}} = \Pi$, and that this is indeed a perfect matching of the possible functions $\Pi : \hat{R} \to \left\{\overline{\mathsf{TF}}, \mathsf{T}\overline{\mathsf{F}}, \mathsf{F}\right\}$, and most importantly that $\Pi[H_1] = \Pi^{\mathsf{c}}[H_2]$ implies $\langle H_1, H_2\rangle \in E$.

Next, observe that for a fixed $\Pi_0 : \hat{R} \to \left\{\overline{\mathsf{TF}}, \mathsf{T}\overline{\mathsf{F}}, \mathsf{F}\right\}$,

$$\Pr_{F_1 \in \mu_q^{R_1}} [\Pi[F_1] = \Pi_0] = \prod_{\hat{\mathsf{f}}:\, \Pi_0(\hat{\mathsf{f}}) = \overline{\mathsf{TF}}} (1-q)^2 \cdot \prod_{\hat{\mathsf{f}}:\, \Pi_0(\hat{\mathsf{f}}) = \mathsf{T}\overline{\mathsf{F}}} q(1-q) \cdot \prod_{\hat{\mathsf{f}}:\, \Pi_0(\hat{\mathsf{f}}) = \mathsf{F}} q$$

Now if $q = p_{max}$, i.e. $(1-q)^2 = q$, we have $\Pr_F[\Pi[F] = \Pi_0] = \Pr_F[\Pi[F] = \Pi_0^{\mathsf{c}}]$. Since $\mu_q(I_1) + \mu_q(I_2) > 1$, there must be a pair Π, Π^{c} such that

$$\{F_1 \in \mathbf{P}(R_1) \mid \Pi[F_1] = \Pi\} \cap I_1 \neq \phi \quad \text{and} \quad \{F_2 \in \mathbf{P}(R_2) \mid \Pi[F_2] = \Pi^{\mathsf{c}}\} \cap I_2 \neq \phi$$

providing the necessary pair of $H_1 \in I_1, H_2 \in I_2$ with $\langle H_1, H_2\rangle \in E$. ■

Lemma 4.10 is thereby proved. ■

The Soundness of the construction (Lemma 4.4) is proven as well. ■

4.3 Tightness

In this section we show our analysis of $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ is tight in two respects. First, we show that for any value of p there is always an independent set $\mathcal{I}$ in $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ whose size is almost $p^\bullet$, regardless of whether or not Φ is satisfiable. Next, we show that if $p > (1-p)^2$ (this happens for $p \geq \frac{3-\sqrt{5}}{2}$), then a large independent set can be formed in $\mathcal{G}_{[p,p^\bullet]}(\Phi)$, again, regardless of the satisfiability of Φ.

Here might be the place to note that our proof of soundness does not mention edges between blocks $B_1, B_2 \in \mathcal{Z}$ for which $|B_1 \cap B_2| < l - 1$ at all. Thus, removing these edges altogether from the graph, would not harm the soundness argument, and would clearly not harm the completeness either.

This is not very surprising since in fact there exists a negligible subset of the vertices $V_0 \subset V$ that covers all of these edges[1]. Thus, a set $I \subset V$ that is independent in the graph with these edges removed, can be translated to $I \setminus V_0$ which is of essentially the same size, and is independent in our graph $\mathcal{G}_{[p,p^\bullet]}(\Phi)$.

In what follows, we ignore such edges, for simpler presentation. Of course, removing V_0 from the independent sets presented here eliminates these additional edges at once.

The 2-intersecting bound. We will exhibit an appropriate choice of maximal 2-intersecting families for almost all of the blocks $\mathcal{Z}$, that constitutes an independent set in $\mathcal{G}_{[p,p^\bullet]}(\Phi)$.

[1] The set V_0 is, in each block, the set of all subsets of R_B for which there is some sub-block $B' \subset B$, $|B'| \leq l - 2$ on which this subset is ”under-represented”, in the sense that less than half of the sub-block-assignments $\mathsf{f} : R_{B'} \to \{\mathsf{T}, \mathsf{F}\}$ have a false extension that falls in the subset.

The reason for the existence of this ”counter-example” is that there is a way to assign almost all $B \in \mathcal{Z}$ with a small special set of block-assignments such that for any two blocks B_1, B_2 such that $|B_1 \cap B_2| = l - 1$, B_1's assignments can be matched to B_2's assignments, so that all but at most one matched pair are consistent. Then, all we need to do is to take, for each block, a 2-intersecting family over this special set of block-assignments. A subset in $\mathcal{I} \cap V[B_1]$, when viewed according to the matching as a subset of the special assignments of B_2, must intersect each subset in $\mathcal{I} \cap V[B_2]$ on at least two special block-assignments, and by the special choice above, at least one of these pairs must be a consistent one.

We exhibit this phenomenon concretely by assigning each block 4 block-assignments, and constructing the ”3 out of 4” 2-intersecting family which is maximal for $p = \frac{3-\sqrt{5}}{2}$.

Let $Z_{red} \cup Z_{green} \cup Z_{blue} \cup Z_{yellow}$ be a partition of Z into roughly equal sizes. For every block $B \in \mathcal{Z}$, define four special block-assignments, $\mathsf{f}^B_{red}, \mathsf{f}^B_{green}, \mathsf{f}^B_{blue}, \mathsf{f}^B_{yellow}$ defined as being true on their color, and false elsewhere, e.g.

$$\forall z \in B, \quad \mathsf{f}^B_{red}(z) \overset{def}{=} \begin{cases} \mathsf{T} & z \in Z_{red} \\ \\ \mathsf{F} & otherwise \end{cases}$$

Of course, not all four are defined for every block, as a block-assignment $\mathsf{f} \in R_B$ must contain at least t T's, and there is a negligible fraction of the blocks $\mathcal{Z}' \subset \mathcal{Z}$ that intersect at least one of $Z_{red} \cup Z_{green} \cup Z_{blue} \cup Z_{yellow}$ with less than t values. Neglecting these, we take for each block, the following set of vertices

$$\mathcal{I}[B] = \left\{ F \in V[B] \;\middle|\; \left| F \cap \left\{ \mathsf{f}^B_{red}, \mathsf{f}^B_{green}, \mathsf{f}^B_{blue}, \mathsf{f}^B_{yellow} \right\} \right| \geq 3 \right\}$$

and let $\mathcal{I} \overset{def}{=} \bigcup_{B \in \mathcal{Z} \setminus \mathcal{Z}'} \mathcal{I}[B]$.

Let $\hat{B} \in Z^{(l-1)}$, and let $B_1 = \hat{B} \cup \{z_1\}$, and $B_2 = \hat{B} \cup \{z_2\}$. Assume $z_1 \in Z_{red}$ (symmetrically for any other color), and observe the following,

1. $\mathsf{f}^{B_1}_{green}, \mathsf{f}^{B_1}_{blue}, \mathsf{f}^{B_1}_{yellow}$ are respectively consistent with $\mathsf{f}^{B_2}_{green}, \mathsf{f}^{B_2}_{blue}, \mathsf{f}^{B_2}_{yellow}$.

2. For any $F_1 \in \mathcal{I}[B_1]$, $\left| F_1 \cap \left\{ \mathsf{f}^{B_1}_{green}, \mathsf{f}^{B_1}_{blue}, \mathsf{f}^{B_1}_{yellow} \right\} \right| \geq 2$, and similarly for $F_2 \in \mathcal{I}[B_2]$, therefore, these vertices are consistent.

Thus, $\mathcal{I}$ is an independent set.

The bound $p < (1-p)^2$. Assume $p > \frac{3-\sqrt{5}}{2}$. We construct an independent set by selecting an arbitrary block assignment for each block, and taking all subsets containing it. By removing a negligible fraction of the vertices (subsets) in each block, we eliminate all edges between blocks.

Consider two blocks $B_1, B_2 \in \mathcal{Z}$, such that $B_1 = \hat{B} \cup \{z_1\}, B_2 = \hat{B} \cup \{z_2\}$. Denote by $\hat{R}$ the set of sub-block assignments for $\hat{B}$ that are restrictions of R_{B_1} and of R_{B_2}, and assume for simplicity that every sub-block assignment in $\hat{R}$ has two extensions (to F and to T) in both R_{B_1} and R_{B_2}.

A random subset $F \in_{\mu_p} \mathbf{P}(R_{B_1})$, has expectedly $p \cdot |R_{B_1}|$ block-assignments. Moreover, there are expectedly $(1-p)^2 \cdot |\hat{R}|$ sub-block-assignments in $\hat{R}$ for which $\mathsf{f}^{(z_1 \leftarrow \mathsf{F})}, \mathsf{f}^{(z_1 \leftarrow \mathsf{T})} \notin F$, and expectedly $p \cdot |\hat{R}|$ sub-block-assignments for which $\mathsf{f}^{(z_1 \leftarrow \mathsf{F})} \in F$.

For two vertices $F_1 \in V[B_1]$ and $F_2 \in V[B_2]$ to be inconsistent, one of them must deviate from the expectation, due to the following. Every $\hat{\mathsf{f}} \in \hat{R}$ for which $\mathsf{f}^{(z_1 \leftarrow \mathsf{F})} \in F_1$ must have both $\mathsf{f}^{(z_2 \leftarrow \mathsf{F})}, \mathsf{f}^{(z_2 \leftarrow \mathsf{T})} \notin F_2$. If both F_1, F_2 are near their expectation, there are roughly $(1-p)^2 \cdot |\hat{R}|$ sub-block-assignments in $\hat{R}$ for which $\mathsf{f}^{(z_2 \leftarrow \mathsf{F})}, \mathsf{f}^{(z_2 \leftarrow \mathsf{T})} \notin F_2$, and since $(1-p)^2 < p$, this is not enough to meet the expected $p \cdot |\hat{R}|$ sub-block-assignments for which $\mathsf{f}^{(z_2 \leftarrow \mathsf{F})} \in F_1$.

Standard Chernoff bounds imply that we need to remove only a tiny fraction of the vertices of each block, so as to eliminate all subsets that deviate from the expectation according to at least one sub-block $\hat{B}$.

Chapter 5

Discussion

Our hardness result for vertex cover relies, in essence, on a new combination of an underlying PCP test-system together with the biased long-code. This PCP system has been custom made for use with the biased long-code, a combination that works due to optimal bounds on the maximal weight of 2-intersecting families.

The Biased Long-Code. The generalization of the long-code, allowing a non-uniform distribution over its bits, opens up new possibilities of utilizing it for obtaining hardness-of-approximation results.

The connection between the study of influences of variables on Boolean functions, and analyzing the long-code, has clarified some of the structure of the long-code, a structure that initially seemed quite complicated. The ideas in this field, and specifically the Friedgut-Lemma, proved to be quite powerful for analyzing the long-code. In particular, the 'decoding' of a p-biased long-code by allowing a slight increase in the value of p in return for low average sensitivity, is actually a reverse view of the well-studied sharp-threshold phenomenon of monotone families.

Standard PCP Terminology and Context. The parallel repetition theorem of [Raz98], in standard PCP terminology says that any NP language L has a 2-prover 1-round interactive proof with certain parameters. Namely, there is a polynomial-time verifier machine V, that upon input I tosses a logarithmic number of coins r and based on that computes queries $q_1(I, r)$ and $q_2(I, r)$ and sends them to provers P_1 and P_2 respectively. The verifier accepts (i.e. declares that $I \in L$) if the provers' answers are consistent. The *completeness* and *soundness* of the proof system are defined as the respective probabilities of accepting an input $I \in L$ and $I \notin L$.

There is a direct translation from this terminology to that of a SAT instance $\langle \Phi, X, Y \rangle$ where the random coin toss r selects a test $\varphi_{x \to y} \in \Phi$, upon which the verifier queries the value of x from prover P_1 and the value of y from prover P_2, and accepts iff $\varphi_{x \to y}(x, y) =$

`true`.

For the maximum independent set problem, the main parameter of interest, is the *free-bit complexity*, as defined by [FK94, BGS98]. This parameter is equal to the average, over the coin-tosses r, of the logarithm$_2$ of the number of possible answers that would cause the verifier to accept. In our terminology, the free bit complexity of our initial SAT instance Φ is simply $\log|R_X|$ because there are $|R_X|$ possible assignments for the X variable and this equals the total number of acceptable answers, since each assignment for the X variable determines exactly one assignment for the Y variable that satisfies $\varphi_{x\to y}$.

comment: Any proof system that has this structure of one prover determining the second prover, can be made into a symmetric proof system by asking both provers the questions of the first prover. This might be interesting later.

A Zero Free-Bit Protocol. It is known that any hardness result for vertex cover can be immediately translated into a PCP protocol with zero free bits as follows. Given input a graph G, the verifier selects a random vertex and accepts if that vertex is outside the vertex-cover and all of its neighbors are inside. This protocol has zero free bits since there is *only one* acceptable answer. Our graph $\mathcal{G}_{[p,p^\bullet]}(\Phi)$ yields such a protocol with completeness $c \approx p$ and soundness $s \approx p^\bullet$. In general, any zero-free-bit protocol with completeness c and soundness s gives a hardness of approximation factor for Vertex-Cover to within $\frac{1-s}{1-c}$.

The Bias Parameter. Currently, the only method for achieving a protocol with few free bits, is via composition: One constructs a PCP protocol e.g. the Raz-verifier or extensions of it, on which one applies a version of the long-code. The trick is to get the right interplay between these two parts so as to achieve the best parameters for the resulting composed protocol.

The failure of the 'naive' construction, namely that of applying the long-code directly onto the FGLSS graph, can be attributed to the fact that in the underlying PCP (parallel repetition) an assignment to the first variable allows for many possible assignments to the second variable. When we apply the p-biased long-code over such a construction, unless p is extremely small, we lose essentially all structure of consistency between variables.

This leads us to define the graph of consistency between answers of the two provers. Let $H_{q_1,q_2} = \langle R_1, R_2, E\rangle$ be a bipartite graph with parts R_1 and R_2 denoting the possible answers of provers 1 and 2 to the respective questions q_1, q_2 (in this paper's terminology: $R_1 = R_X$, and $R_2 = R_Y$). Connect $a_1 \in R_1$ and $a_2 \in R_2$ by an edge if (a_1, a_2) will cause the verifier to reject. Note that H_{q_1,q_2} is the subgraph consisting of two layers in the FGLSS graph. Applying the long-code over H_{q_1,q_2} simply means to take the bipartite graph $\mathcal{LC}(H_{q_1,q_2}) = \langle \mathbf{P}(R_1), \mathbf{P}(R_2), E_{LC}\rangle$ whose parts are $\mathbf{P}(R_1)$ and $\mathbf{P}(R_2)$ and put an edge between $F_1 \in \mathbf{P}(R_1)$ and $F_2 \in \mathbf{P}(R_2)$ if $F_1 \times F_2 \subseteq E$.

Intuitively, the sparser we can make H_{q_1,q_2}, the larger we can take p in the p-biased–long-code to be (our ultimate goal being $p = \frac{1}{2}$).

Call a vertex set *nearly*(p)-independent in $\mathcal{LC}(H_{q_1,q_2})$, if one can remove from it a set whose size according to μ_p is exponentially small in $|R_1|, |R_2|$ to make it an independent set in $\mathcal{LC}(H_{q_1,q_2})$. Define the p-threshold-family $\mathcal{F}^{\geq p}(R) = \{F \subset R \mid |F| \geq p \cdot |R|\}$, and let

$$\text{bias}(H_{q_1,q_2}) \overset{def}{=} \sup_p \left(\mathcal{F}^{\geq p}(R_1) \cup \mathcal{F}^{\geq p}(R_2) \text{ is } \textit{not} \text{ nearly}(p)\text{-independent in } \mathcal{LC}(H_{q_1,q_2})\right).$$

Finally, we define the bias parameter of the proof system to be the maximum, over all pairs q_1, q_2, of bias(H_{q_1,q_2}). Notice that if the average co-degree of H_{q_1,q_2} is large, we have no hope of getting a large bias parameter. In fact, it would have been nice, but appears non-trivial, to be able to formulate a connection between the average or the minimum co-degree of H and an upper bound on the bias of H (the expansion of the graph with respect to sets of relative size p can be used to lower-bound the bias parameter, but this connection is not tight).

Let us now inspect our construction in view of this new parameter. Indeed, our graph $\mathcal{G}_{\mathcal{Z}}[\Phi]$ can be viewed as a 1-round 2-prover proof-system with bias parameter $p = \frac{3-\sqrt{5}}{2}$ as follows. Let the question to the first prover be a block $B \in \mathcal{Z}$, and a possible answer a block-assignment in R_B. Let the question to the second prover be symmetrically another block B', intersecting with the previous one on an $l-1$-sub-block $B \cap B' = \hat{B}$, and again an answer would be a block-assignment in $R_{B'}$. The verifier will accept if the answers agree on the common $l-1$ sub-block. Obviously the graph $H_{B,B'}$ is the subgraph of $\mathcal{G}_{\mathcal{Z}}[\Phi]$ induced by two blocks.

Interestingly, although this construction has almost perfect completeness (i.e. a good input is accepted with probability $1-\epsilon$) we cannot prove that the soundness probability is low. However, for establishing a lower-bound for the vertex cover problem, a weaker gap suffices. Indeed, we prove that it is NP hard to distinguish between the following:

- There is a prover-strategy causing the verifier to accept with probability $1-\epsilon$.
- For any prover strategy, any ϵ fraction of queries (i.e. blocks) contain at least one inconsistent pair.

The 2-Intersecting Limitation. When we apply the biased long-code to the above PCP protocol, the composition limits our bias to $p \leq \frac{3-\sqrt{5}}{2}$. A further limitation of our technique results in a gap not between $1-p$ and $1-\epsilon$ but rather between $1-p$ and $1-p^{\bullet}$. The reason for this is that the underlying PCP system (corresponding to $\mathcal{G}_{\mathcal{Z}}[\Phi]$) does not allow for "multi-assignments". Regardless of the original input, we can exhibit a choice of two answers per query, such that for any pair of blocks there is always a selection of one from each that make an acceptable answer. Had this not been the case, a stronger vertex-cover result would follow. Let us state a conjecture that would imply hardness of approximating vertex cover to within $2-\epsilon$:

Conjecture 5.1 *For every $\epsilon, h > 0$ there exists an $R > 1$, such that if Φ is a SAT instance with variables X whose range is R, such that each test $\varphi \in \Phi$ depends on two variables and the bias parameter is $\frac{1}{2}$, then it is NP-hard to distinguish between*

- *There is an assignment $A : X \to R$ satisfying $1 - \epsilon$ of Φ.*
- *For every multi-assignment $A : X \to \binom{R}{h}$ there is a pair of variables whose multi-assignments are pairwise inconsistent.*

Moreover, as explained above, such a conjecture would also immediately imply the existence of a PCP with 0 free-bits, soundness ϵ and completeness $\frac{1}{2} - \epsilon$.

Other problems. Analyzing existing results for other approximation problems with tools such as the Friedgut Lemma, may be fruitful. Two obvious candidates for this are the Maximum-Clique, and the problem of approximating the Chromatic Number of a 3-colorable graph.

The best known algorithm for coloring such a graph, uses n^{const} colors [KMS98]. On the other hand, it is known to be NP-hard to color such a graph with 5 colors [KLS00]. Narrowing this gap is one of the most interesting remaining open questions in this field.

Such a hardness result can be obtained, for example, if one exhibits a graph G for which it is NP-hard to distinguish between the case G is 3-colorable and the case where the largest independent-set of G is small.

One of the obstacles preventing the adaptation of our construction to this setting, is the fact that our construction has imperfect completeness. Even in case of a satisfiable instance, almost but not all of $\mathcal{Z}$ participates in the independent set. This prevents our graph from being 3-colorable, in case we start with a satisfying assignment.

Part II

The Closest Vector Problem

Chapter 6

Introduction

An n-dimensional lattice $L = L(v_1, .., v_n)$, for linearly independent vectors $v_1, .., v_n \in R^k$ is the additive group generated by the vectors, i.e. the set $L = \{\sum a_i v_i \mid a_i \in \mathbb{Z}\}$. Given L and an arbitrary vector y, the Closest Vector Problem (CVP) is to find a vector in L closest to y in a certain norm. The Shortest Vector Problem (SVP) is a homogeneous analog of CVP, and is defined to be the problem of finding the shortest non-zero vector in L.

These lattice problems have been introduced in the 19th century, and have been studied since. Minkowsky and Dirichlet tried, with little success, to come up with approximation algorithms for these problems. It was much later that the lattice reduction algorithm was presented by Lenstra, Lenstra and Lovász [LLL82] , achieving a polynomial-time algorithm approximating the Shortest Lattice Vector to within the exponential factor $2^{n/2}$, where n is the dimension of the lattice. Babai [Bab86] applied LLL's methods to present an algorithm that approximates CVP to within a similar factor. Schnorr [Sch85] improved on LLL's technique, reducing the factor of approximation to $(1+\epsilon)^n$, for both CVP and SVP, where the polynomial running time depends on $\frac{1}{\epsilon}$ in the exponent. These positive approximation results are quite weak, achieving only extremely large (exponential) factors. The question naturally arises: What are the factors of approximation to within which these problems can be approximated in polynomial time?

Interest in lattice problems has been renewed due to a result of Ajtai [Ajt96], showing a reduction, from a version of SVP, to the *average-case* of the same problem.

CVP was shown to be NP-hard for any l_p norm in [vEB81], where it was also conjectured that SVP is NP-hard. Arora et al. [ABSS93] utilized the PCP characterization of NP to show that CVP is NP-hard to approximate to within any constant, and quasi-NP-hard to approximate to within $2^{(\log n)^{1-\epsilon}}$ for any constant $\epsilon > 0$ (i.e. an approximation algorithm for such factors would imply $NP \subseteq \text{DTIME}(2^{polylogn})$.

As for SVP, it is NP-hard to approximate to within [Ajt98, Mic98] some constant factor (see also [CN98]). The proof in [Mic98] relies on the PCP characterization of NP and is carried out via a reduction from gap-CVP (shown NP-hard for any constant gap in [ABSS93]). Using gap-CVP allows, in addition to the significant improvement in the hardness-of-approximation factor, a major simplification of the main technical lemma from [Ajt98]. Better hardness results for gap-CVP may result in improved approximation hardness results for SVP.

So far there is still a huge gap between the positive results, showing approximations for these problems with exponential factors, and the above hardness results. Nevertheless, some other results provide a discouraging indication for improving the hardness result beyond a certain factor. [LLS90] showed that approximating CVP to within $n^{1.5}$ is in co-NP, and later [GG98] showed that approximating both SVP and CVP to within $\sqrt{n}$ is in NP$\cap$ co-AM. Hence it is unlikely for any of these problems to be NP-hard.

The strongest NP-hardness result likely to be true for these problems, hence, is that they are NP-hard to approximate to within a constant power of the dimension.

Our Results. We improve on [ABSS93] in two ways. First, we go beyond the factor of $2^{(\log n)^{1-\epsilon}}$ for any constant $\epsilon > 0$, which was the previous hardness-of-approximation factor known for CVP. Instead, we achieve a factor of $2^{\frac{\log n}{\log \log n}} = n^{\frac{1}{\log \log n}}$. Furthermore, we show approximating CVP is *NP-hard* for these large factors, compared to the previously known *quasi* NP-hardness.

The known PCP characterizations of NP seem inadequate in order to show hardness of approximating CVP to within large factors. The proof of [ABSS93] utilizes amplification techniques, in which the dimension of the instance grows faster than the factor for which approximation hardness is obtained. It is therefore unlikely that using this technique, even if allowing a super-polynomial blow-up, one can obtain such strong results. It seems that with this method it will always be the case that the factor for which hardness of approximation is proven never reaches beyond the barrier of $2^{(\log n)^{1-\epsilon}}$ for any constant $\epsilon > 0$.

We introduce a new NP-hard gap-problem, Super-SAT (S-SAT for short), that we use to prove our result. The S-SAT problem is a gap version of SAT, minimizing a new, appropriately defined, objective function. Although the S-SAT characterization differs from the PCP characterizations, its proof relies on similar techniques.

Let SAT[F] be the following problem: An instance of SAT[F] is a set of local-constraints (Boolean functions) called *tests*, on variables from a common set, each variable ranging over a finite set F. Each test is represented by a list of assignments for its variables, which are said to satisfy the test. The goal is to attach to each test one of the assignments that satisfies it, such that consistency is maintained among the assignments, that is, each variable is given the same value by the assignments of all tests that depend on it. If this is possible, the instance is accepted, and otherwise it is rejected.

Our gap version of this problem, S-SAT, is as follows: S-SAT is the same as SAT[F] except not all non-satisfiable instances must be rejected. We generalize the notion of assignment to that of super-assignment – formal linear combinations of assignments with integer coefficients – and modify the acceptance condition accordingly: Previously accepted instances must still be accepted. An instance must be rejected only if there is no super-assignment to the tests, whose norm (see Definition 7.2) is smaller than g, and which is "everywhere consistent" (in a sense similar to that described above). If the instance is somewhere in-between (i.e. minimizing the norm of its consistent super-assignments gives a value greater than 1 but less than g), then that instance is not necessarily rejected (any outcome is ok).

We show (Theorem 7.1) that solving this problem is NP-hard for $g = n^{1/\log\log n}$ (n denotes, as usual, the size of the instance). We then reduce this problem to CVP, preserving the approximation factor. Improving the hardness of approximation factor of S-SAT to a constant power of n, namely where $g = n^{\epsilon}$ for some constant ϵ (Conjecture 7.2), would directly imply CVP to be hard to approximate to within a constant power of the dimension.

For simplicity, our proof works with l_1 norm, however it can be extended to l_p norm for any $1 < p < \infty$ as shown in Section 11.3.

Outline. We begin, in Chapter 7, by presenting the new NP-hard gap-problem, S-SAT. We first formally define S-SAT and then state Theorem 7.1 asserting it is NP-hard to approximate to within large factors of approximation ($n^{1/\log\log n}$). Chapter 8 gives some definitions and techniques which are the basis of the construction. The NP-hardness of S-SAT, that is the most technical part of this work, is established in two parts. In Chapter 9, we describe the reduction from a low error-probability PCP characterization of NP, to S-SAT. We proceed to prove the correctness of the reduction (Theorem 7.1) in Chapter 10. In Finally, in Chapter 11 we show a simple reduction from S-SAT to CVP.

Chapter 7

Super-SAT - S-SAT

In this chapter we introduce a new NP-hard problem, S-SAT. Let us begin by defining SAT[$\mathcal{F}$], which is actually SAT over non-Boolean variables, presented from a different point of view. An instance of SAT[$\mathcal{F}$]

$$I = \langle \Psi = \{\psi_1, .., \psi_n\}, V = \{v_1, .., v_m\}, \{\mathcal{R}_{\psi_1}, .., \mathcal{R}_{\psi_n}\} \rangle$$

is a set Ψ of *tests* (Boolean functions) over a common set V of variables that take values in a field $\mathcal{F}$. In what follows $|\mathcal{F}|$, m, and $|\mathcal{R}_{\psi_i}|$ will always be bounded by a polynomial in $n = |\Psi|$. Each test $\psi \in \Psi$ has associated with it a list $\mathcal{R}_\psi$ of assignments to its variables, called the *satisfying assignments* or the *range* of the test ψ. Having both ψ and $\mathcal{R}_\psi$ is convenient yet somewhat redundant since the list $\mathcal{R}_\psi$ actually specifies all there is to know about the test ψ.

An *assignment* for an instance maps to each *test*, a satisfying assignment from its range. An instance is accepted iff there is an assignment to the tests that is everywhere consistent, that is, each variable is given the same value by the assignments to all tests that depend on it. It is easy to see that SAT[$\mathcal{F}$] is NP-complete.

S-SAT is a gap variant of this problem, obtained by setting a new measure on the non-satisfiability of an instance. While in PCP we measured the fraction of tests, satisfiable by a single assignment, in S-SAT we will define a measure of a different nature - we will introduce a notion of super-assignments to the tests, that is, formal linear combinations of assignments. We will then measure the 'length' of a super-assignment, and ask how 'short' it may get while maintaining 'consistency'.

Definition 7.1 (Super-Assignment to Tests) *A super-assignment is a function S mapping to each $\psi \in \Psi$ a value from $\mathbb{Z}^{\mathcal{R}_\psi}$. $S(\psi)$ is a vector of integer coefficients, one for each value $r \in \mathcal{R}_\psi$. Denote by $S(\psi)[r]$ the r^{th} coordinate of $S(\psi)$.*

If $S(\psi)[r] \neq 0$, we say that the value r *appears* in $S(\psi)$. A *natural super-assignment* assigns each $\psi \in \Psi$ a unit vector $e_i \in \mathbb{Z}^{\mathcal{R}_\psi}$ with a 1 in a single coordinate i corresponding to an assignment for that test in the usual sense (i.e. an assignment which maps $r \in \mathcal{R}_\psi$ to ψ corresponds to the natural super-assignment $S(\psi)$ such that $S(\psi)[r] = 1$ and $S(\psi)[r'] = 0$ for all $r' \neq r$). We use the average over the l_1 norms of the vectors $S(\psi)$, $\|S(\psi)\|$, to measure the closeness of S to a natural super-assignment,

Definition 7.2 (Norm of a Super-Assignment) *The norm of a super-assignment S is the average norm of its individual assignments $\|S\| = \frac{1}{|\Psi|}\sum_{\psi\in\Psi}\|S(\psi)\|$, where $\|S(\psi)\|$ denotes the l_1 norm of the vector $S(\psi)$.*

The norm of a natural super-assignment is 1. The gap of S-SAT will be formulated in terms of the norm of the minimal super-assignment that maintains consistency. A natural assignment $r \in \mathcal{R}_\psi$ to a test ψ induces an assignment to each variable x, denoted $r|_x$. In the SAT[$\mathcal{F}$] problem an assignment is called consistent if for every pair of tests with a common variable, the assignments to the tests, restricted to the variable, are equal. We extend this notion of consistency to super-assignments by defining the projection of a super-assignment $S(\psi)$ onto each of ψ's variables. Consistency between tests will amount to equality of projections on common variables.

Definition 7.3 (Projection) *Let $S : \Psi \to \bigcup_\psi \mathbb{Z}^{\mathcal{R}_\psi}$ be a super-assignment to the tests. We define the projection of $S(\psi)$ on a variable x of ψ, $\pi_x(S(\psi)) \in \mathbb{Z}^{\mathcal{F}}$, as follows:*

$$\forall a \in \mathcal{F}: \qquad \pi_x(S(\psi))[a] \overset{def}{=} \sum_{r\in\mathcal{R}_\psi,\, r|_x = a} S(\psi)[r]$$

Namely, we partition the assignments in $\mathcal{R}_\psi$ according to their value $a \in \mathcal{F}$ on the variable x (we associate with $a \in \mathcal{F}$ all assignments $r \in \mathcal{R}_\psi$ for which $r|_x = a$). For each value $a \in \mathcal{F}$, we then add the coefficients $S(\psi)[r]$ of the assignments associated with it, and this is the value of the coefficient $\pi_x(S(\psi))[a]$.

We shall now proceed to define the notion of consistency between tests. If the projections of two tests on each common variable x are equal (in other words, they both give x the same super-assignment), we say that the super-assignments of the tests are consistent.

Definition 7.4 (Consistency) *Let S be a super-assignment to the tests in Ψ. S is consistent if for every pair of tests ψ_i and ψ_j with a common variable x,*

$$\pi_x(S(\psi_i)) = \pi_x(S(\psi_j))$$

S is said to be *non-trivial* if every variable $x \in V$ there is at least one test $\psi \in \Psi$ that isn't 'cancelled' on x: $\pi_x(S(\psi)) \neq \vec{0}$. For a variable x we think of all the values $a \in \mathcal{F}$ receiving non-zero coefficients in $\pi_x(S(\psi))$ (i.e. values for which $\pi_x(S(\psi))[a] \neq 0$)) as being simultaneously 'assigned' to x by ψ. The non-triviality requirement means that each variable must be assigned at least one value.
We can now define the S-SAT problem.

Definition 7.5 (g-S-SAT) *An instance of* S-SAT *with parameter* g

$$I = \langle \Psi = \{\psi_1, .., \psi_n\}, V = \{v_1, .., v_m\}, \{\mathcal{R}_{\psi_1}, .., \mathcal{R}_{\psi_n}\}\rangle$$

consists of a set Ψ *of* tests *over a common set* V *of variables that take values in a field* $\mathcal{F}$*. The parameters* m *and* $|\mathcal{F}|$ *and* $|\mathcal{R}_\psi|$ *are always bounded by some polynomial in* n*. Each test* $\psi \in \Psi$ *has associated with it a list* $\mathcal{R}_\psi$ *of assignments to its variables, called the* satisfying assignments *or the* range *of the test* ψ*. The problem is to distinguish between the following two cases,*

Yes: There is a consistent natural super-assignment for Ψ*.*

No: Every non-trivial consistent super-assignment for Ψ *has norm* $> g$*.*

Theorem 7.1 (S-SAT **Theorem**) *There is some constant* $c > 0$*, such that* S-SAT *is NP-hard for* $g = n^{c/\log\log n}$*.*

The S-SAT theorem (Theorem 7.1) can be viewed as an extension of Cook's theorem [Coo71, Lev73] in the following way. An algorithm solving S-SAT is required to accept if the test system is satisfiable. However, the algorithm is allowed to accept non-satisfiable instances that have a consistent super-assignment of norm $\leq g$. It must only reject when every consistent super-assignment for Ψ has norm $> g$. We are, in fact, adding slackness between the acceptance and rejection cases.

We suggest a stronger conjecture which, if true, would imply that CVP is NP-hard to approximate to within a *constant power* of the lattice-dimension.

Conjecture 7.2 S-SAT *is NP-hard for* $g = n^c$ *for some constant* $c > 0$*.*

It is unlikely that the conjecture remain true for $c \geq \frac{1}{2}$ due to the result of [GG98] showing that approximating CVP to within $\sqrt{n}$ is in NP$\cap$ co-AM. Our reduction from S-SAT to CVP is linear, and hence it follows that approximating S-SAT to within $\sqrt{n}$ is in NP$\cap$ co-AM as well.

Chapter 8

Tools and Definitions

8.1 Preliminaries

Let $\mathcal{F}$ denote a finite field $\mathcal{F} = \mathbb{Z}_p$ for some prime number $p > 1$.

Definition 8.1 (Low Degree Function) *A function $f : \mathcal{F}^d \to \mathcal{F}$ is said to have degree r if its values are the point evaluation of a polynomial on $\mathcal{F}^d$ with degree $\leq r$ in each variable. In this case we say that f is an $[r, d]$-LDF, or $f \in LDF_{r,d}$.*

Sometimes we omit the parameters and refer simply to an LDF. The *total degree* of a function is the total degree of the corresponding polynomial, i.e. the maximum over its monomials, of the sum of degrees of each variable in the monomial. Every $[r, d]$-LDF has total degree at most rd.

For an LDF $P : \mathcal{F}^d \to \mathcal{F}$, we define its restriction and re-parameterization $P|_{\mathcal{C}} : \mathcal{F}^D \to \mathcal{F}$ to the D dimensional cube (affine subspace) $\mathcal{C} = \bar{x}_0 + \mathrm{span}(\bar{x}_1, .., \bar{x}_D)$ (where $\bar{x}_0, .., \bar{x}_D \in \mathcal{F}^d$), in the natural way. Namely,

$$\forall (t_1, .., t_D) \in \mathcal{F}^D, \qquad P_{\mathcal{C}}(t_1, .., t_D) = P(\bar{x}_0 + \sum_{i=1}^{D} t_i \bar{x}_i)$$

Observe that the total degree of $P|_{\mathcal{C}}$ is at most that of P, namely $\leq rd$.

Definition 8.2 (Low Degree Extension) *Let $m, d > 0$ be natural numbers, and let $\mathcal{H} \subset \mathcal{F}$ such that $|\mathcal{H}|^d = m$. A vector $(a_0, .., a_{m-1}) \in \mathcal{F}^m$ can be naturally identified with a function $A \colon \mathcal{H}^d \to \mathcal{F}$ by looking at points in $\mathcal{H}^d$ as representing numbers in base $|\mathcal{H}|$.*

Let $\hat{A} : \mathcal{F}^d \to \mathcal{F}$ be defined by

$$\hat{A}(x_1, .., x_d) = \sum_{(h_1,..,h_d) \in \mathcal{H}^d} \prod_{\substack{i \,\in\, \mathcal{H} \\ i \,\neq\, h_1}} \frac{(x_1 - i)}{(h_1 - i)} \cdot \prod_{\substack{i \,\in\, \mathcal{H} \\ i \,\neq\, h_2}} \frac{(x_2 - i)}{(h_2 - i)} \cdots \prod_{\substack{i \,\in\, \mathcal{H} \\ i \,\neq\, h_d}} \frac{(x_d - i)}{(h_d - i)} \cdot A(h_1, .., h_d)$$

$\hat{A}$ is a $(|\mathcal{H}|-1, d)$-LDF called the $|\mathcal{H}|-1$ degree extension of A in $\mathcal{F}$.

Let $V = \{v_0, .., v_{m-1}\}$ be a set of variables, and identify every assignment $A : V \to \mathcal{F}$ with the vector $(a_0, .., a_{m-1}) \in \mathcal{F}^m$ where $a_i = A(v_i)$. One can extend A to a larger set of variables $\hat{V} \supset V$ via the low-degree-extension of $(a_0, .., a_{m-1})$. Namely, we identify the variables V with the points in $\mathcal{H}^d$, and add new variables for the rest of the points in $\mathcal{F}^d$. The new set of variables $\hat{V}$ correspond each to a point in $\mathcal{F}^d$. $\hat{A}$ is thus viewed as an assignment to $\hat{V} \supset V$ that (1) extends A, and (2) is a point-evaluation of an $[|\mathcal{H}|-1, d]$-LDF.

Similar to the definition of super-assignments, we define a *super-LDF* $\mathcal{G}$: LDF $\to \mathbb{Z}$ to be a formal integer linear combination of LDFs, and denote by $\mathcal{G}[P]$, the integer coefficient assigned to the LDF P. We say that the LDF P *appears* in $\mathcal{G}$ iff $\mathcal{G}[P] \neq 0$. This definition arises naturally from the fact that the tests in our final construction will range over LDFs. We further define the *norm* of a super-LDF to be the norm of the corresponding coefficient vector (same as with super-assignments). We say that a super-LDF has total degree r if every LDF appearing in it has total degree $\leq r$.

Given a super-$[r, d]$-LDF $\mathcal{G}$, we define its restriction $\pi_{\mathcal{C}}(\mathcal{G})$ to a D-dimensional cube $\mathcal{C}$, (which is a super-LDF of dimension D and degree rd) in the natural way. Namely,

$$\forall P \in \mathrm{LDF}_{rd,D} \qquad \pi_{\mathcal{C}}(\mathcal{G})[P] \overset{def}{=} \sum_{Q \in \mathrm{LDF}_{r,d},\, Q|_{\mathcal{C}} = P} \mathcal{G}[Q]$$

We say that a point x is ambiguous for a super-LDF $\mathcal{G}$ if there are two LDFs appearing in $\mathcal{G}$, that agree on x. The following (simple) property of super-LDFs will be very important.

Proposition 8.1 (Low Ambiguity) *Let $\mathcal{G}$ be an $[r, d]$-super-LDF of norm $\leq g$. The fraction of ambiguous points for $\mathcal{G}$ is $\leq \mathrm{amb}(r, d, g) \overset{def}{=} \binom{g}{2}\frac{rd}{|\mathcal{F}|}$.*

Proof: Two distinct $[r, d]$-LDFs agree on at most $\frac{rd}{|\mathcal{F}|}$ of their points. At most g LDFs appear in any super-LDF of norm $\leq g$, and so there are no more than $\binom{g}{2}$ pairs. ■

Two LDFs can coincide on only a small fraction of cubes,

Proposition 8.2 *Let P, Q be two $[r, d]$-LDFs. The fraction of cubes $\mathcal{C}$ (affine subspaces of dimension $D < d$) on which $P|_{\mathcal{C}} = Q|_{\mathcal{C}}$ is $\leq \frac{rd}{|\mathcal{F}|}$.*

This follows from the fact that two distinct $[r, d]$-LDFs agree on at most $\frac{rd}{|\mathcal{F}|}$ of their domain, and by the fact that selecting a random point in a random cube gives a uniform distribution on the entire domain, which implies that the restriction of an LDF to a random cube, is even less likely to avoid all points for which $P(x) \neq Q(x)$.

8.2 Embedding Extension

An important technique utilized herein is adapted from [DFK$^+$99], and shows how to represent an LDF over a low-dimensional domain $\mathcal{C} = \mathcal{F}^t$ by a lower-degree LDF over a domain of higher dimension $\mathcal{D} = \mathcal{F}^{kt}$. The points in the domain $\mathcal{C}$ are embedded in the domain $\mathcal{D}$ by taking each 'axis' in $\mathcal{F}^t$ and replacing it by k new ones (thus the extended domain $\mathcal{F}^{kt}$ has dimension $k \cdot t$) so that an LDF of degree r (in each variable) on the original domain $\mathcal{F}^t$ is transformed to an LDF of degree $\sqrt[k]{r}$ (in each variable) on the extended domain $\mathcal{F}^{kt}$.

Definition 8.3 (embedding extension) *Let $b \geq 2$, $k > 1$ and t be natural numbers. We define the* embedding extension mapping $E_b : \mathcal{F}^t \to \mathcal{F}^{t \cdot k}$ *as follows. E_b maps any point $x = (\xi_1, .., \xi_t) \in \mathcal{F}^t$ to $y \in \mathcal{F}^{t \cdot k}$, $y = E_b(x) = (\eta_1, .., \eta_{t \cdot k})$ by*

$$E_b(\xi_1, .., \xi_t) \overset{def}{=} \left(\xi_1, (\xi_1)^b, (\xi_1)^{b^2}, .., (\xi_1)^{b^{k-1}}, \ldots, \xi_t, (\xi_t)^b, (\xi_t)^{b^2}, .., (\xi_t)^{b^{k-1}}\right)$$

Hence $E_b(\mathcal{F}^t) \subset \mathcal{F}^{kt}$ is a manifold (multi-dimensional curve) in $\mathcal{F}^{kt}$. Each of $\mathcal{F}^{kt}$'s axes corresponds to some preset power of an axis of $\mathcal{F}^t$, and $E_b(\mathcal{F}^t)$ consists of exactly the points in which those axes indeed match.

The following proposition shows that any LDF on $\mathcal{F}^t$ can be represented by an LDF on $\mathcal{F}^{t \cdot k}$ with significantly lower degree:

Proposition 8.3 *Let $f : \mathcal{F}^t \to \mathcal{F}$ be a $[b^k - 1, t]$-LDF, for integers $t > 0, b > 1, k > 1$. There is a $[b - 1, t \cdot k]$-LDF $f_{ext} : \mathcal{F}^{t \cdot k} \to \mathcal{F}$ such that*

$$\forall x \in \mathcal{F}^t : \quad f(x) = f_{ext}(E_b(x))$$

Proof: We rewrite f as an LDF $f_{ext} : \mathcal{F}^{t \cdot k} \to \mathcal{F}$ by replacing each power $(\xi_i)^p$ of ξ_i,

$$0 < i \leq t \quad 0 < p < b^k \qquad (\xi_i)^p \longrightarrow (\eta_{i,0})^{\beta_0} \cdot (\eta_{i,1})^{\beta_1} \cdots (\eta_{i,k-1})^{\beta_{k-1}}$$

where $\langle \beta_0 \beta_1 ... \beta_{k-1} \rangle$ is the base b representation of p, and we 're-index' $\eta_{i,j} \overset{def}{=} \eta_{(i-1)k+j+1}$. The degree in each variable of f_{ext} is $b - 1$, and the dimension is $t \log_b b^k = t \cdot k$. The restriction of f_{ext} to the manifold $E_b(\mathcal{F}^t)$, will give f, as seen from substituting the manifold equations $\eta_{i,j} = (\eta_{i,0})^{b^j}$ into each of the monomials). ■

Note that an arbitrary $[r, tk]$-LDF f on the larger domain $\mathcal{F}^{t \cdot k}$ can be viewed, when restricted to the manifold, as a $[\tilde{r}, t]$-LDF $\tilde{f}$ with $\tilde{r} = r \cdot (1 + b + b^2 + \ldots + b^{k-1}) \leq r \cdot (b^k - 1)$. This LDF is the re-parameterization of the LDF obtained by substituting in the manifold equations. Note that if the total degree of f is s, then the total degree of $\tilde{f}$ is $\leq s \cdot b^{k-1}$.

Chapter 9

Reducing PCP to S-SAT

In this chapter, we present a reduction from a low error-probability PCP characterization of NP, to S-SAT. Starting with a PCP instance, we show how to construct an instance of S-SAT. The correctness of the reduction is proven in the next chapter.

Let $\Phi = \{\varphi_1, .., \varphi_n\}$ be a system of *tests* over Boolean variables $V_\Phi = \{v_1, .., v_m\}$, (assume $m = n^c$ for some constant $c > 0$) such that each test depends on $D = O(1)$ variables. The following theorem is a direct corollary of [AS92, ALM$^+$92]:

Theorem 9.1 *It is NP-hard to distinguish between the following two cases:*

Yes: There is an assignment to V_Φ such that all $\varphi_1, ..., \varphi_n$ are satisfied.

No: No assignment can satisfy more than $1/2$ of the tests in Φ.

Starting from Φ, we will construct an S-SAT test-system Ψ over variables $V_\Psi \supset V_\Phi$. Our new variables V_Ψ will range over a larger, non-Boolean, range, namely a field $\mathcal{F}$. An assignment to V_Ψ can be interpreted as an assignment to V_Φ by identifying the value $0 \in \mathcal{F}$ with the Boolean value `true` and every non-zero value $a \in \mathcal{F}$ with the Boolean value `false`.

9.1 Constructing the CR-Forest

We construct Ψ from Φ by replacing each $\varphi \in \Phi$ with a set of new tests ψ. These tests essentially test that φ is satisfied, and that some set of variables (that encode φ's variables) are an LDF. The construction relies on strong 'error-correcting' properties of LDFs (in a similar manner to proofs of PCP theorems) to eventually 'decode' any consistent low-norm super-assignment for Ψ into a satisfying assignment for the original test-system Φ. The idea is to embed Φ's variables into a geometric domain and then recursively encode this domain by multiple new domains, adding new variables along the way.

We describe the construction via an underlying tree structure, one tree per test $\varphi \in \Phi$. Each node in the tree is associated with a set of variables such that the variables of all of the offspring of a node encode that node's variables. For each leaf of the tree, Ψ will have one test that depends on the variables associated with that leaf.

The key to the construction lies in understanding how the variables associated with different nodes relate to each other This is described in Section 9.2. The variables of the root node contain φ's variables, plus some additional ones that together represent the points of a domain $\mathcal{F}^{d_0}$. In fact, every node in the tree will associated with a domain $\mathcal{F}^d$, and each offspring of that node will be associated with a cube $\mathcal{C} \subset \mathcal{F}^d$ in that domain. This is roughly how the points of the parent domain are distributed among its offspring. The variables of each offspring will consist of some of the parent's variables but also some new "extension" variables, together corresponding to points in a new domain $\mathcal{F}^d$ where the parent "cube" variables are mapped via the embedding extension mapping into the new domain.

The idea is that a consistent super-assignment to the tests of Ψ, essentially assigning a super-LDF to each leaf node, can be inductively decoded into super-LDFs on domains of nodes residing higher up in the tree, reaching all the way up to the root. For this decoding to work, certain points in a domain, containing more 'information' than others, need to have a larger proportion of offspring representing them. This is established (in Section 9.3) by defining for each domain a set of 'distinguished points'. Then, a mechanism of labels serves to obtain the correct proportion of offspring encoding the distinguished and the non-distinguished points.

Let us begin by defining the composition-recursion forest (CR-forest), which holds the underlying structure of Ψ.

Let $\mathcal{F}$ be a field of size $|\mathcal{F}| = |V_\Phi|^{\Theta(1)/\log\log n} = n^{c_1/\log\log n}$ for some constant $c_1 > c$ (recall we denoted $|V_\Phi| = n^c$). Let $d_0 = \lceil \log\log n \rceil$, recall that D denotes the number of variables each test in Φ depends on, and set $d = 4D + 8$. Let $L = \lceil c_2 \log\log n \rceil$, (the constant $c_2 > 0$ will be specified later).

Let $B(\mathcal{F}^d, t_1, t_2)$ denote the number of different affine-subspaces of dimension t_1 (in a domain $\mathcal{F}^d$) that contain a certain affine subspace of dimension $0 \le t_2 \le t_1$. It is easy to see that $B(\mathcal{F}^d, t_1, t_2) \le |\mathcal{F}|^{d(t_1 - t_2)}$.

Definition 9.1 ($\mathbf{F}_n(\Phi)$) *The composition-recursion forest (CR-forest)* $\mathbf{F}_n(\Phi) = \{\mathbf{T}_\varphi \mid \varphi \in \Phi\}$ *is a set containing one depth-L tree $\mathbf{T}_\varphi$ for every test $\varphi \in \Phi$. The root node (level-0) of $\mathbf{T}_\varphi$ has $B(\mathcal{F}^{d_0}, D+2, D-1) = n^{O(1)}$ offspring, and all nodes in levels $i = 1, \ldots, L-1$ have $2\,|\mathcal{F}|^{D+2} \cdot B(\mathcal{F}^{4D+8}, D+2, 0) = |\mathcal{F}|^{O(1)}$ offspring. Note that although the forest $\mathbf{F}_n(\Phi)$ depends on many parameters (L, D, d_0) which can all be derived from Φ, we single out the parameter n according to which the size of the generated instance is measured.*

The forest $\mathbf{F}_n(\Phi)$ will be the base upon which Ψ's variables and tests will be defined as follows. With each node $v \in \mathbf{T}_\varphi$ $(\varphi \in \Phi)$, we associate a distinct geometric domain, denoted $\mathbf{dom}_v$. For the root $root_\varphi$ of every tree, $\mathbf{dom}_{root_\varphi} \overset{def}{=} \mathcal{F}^{d_0}$, while for non-root nodes

v, $\mathbf{dom}_v \overset{def}{=} \mathcal{F}^d$. For a node v, we associate with each point in $\mathbf{dom}_v$ a distinct variable from V_Ψ, by defining an injection $\mathbf{var}_v : \mathbf{dom}_v \to V_\Psi$. Points from domains of distinct nodes may be mapped to the same variable. In particular, the variables that φ depends on will belong to $\mathbf{var}_v(\mathbf{dom}_v)$ for many of the leaves in the tree $\mathbf{T}_\varphi$.

We can already at this point define the tests of Ψ,

Definition 9.2 (tests) *Ψ will have one test ψ_v for each leaf v in the forest. ψ_v will depend on the variables in $\mathbf{var}_v(\mathbf{dom}_v)$. An assignment A for ψ_v's variables is considered satisfying if and only if the following two conditions hold:*

1. *A is an $[r_L, d]$-LDF on $\mathbf{var}_v(\mathbf{dom}_v)$ (where $r_L \le 2(D+2) = O(1)$ will be defined below).*
2. *If $v \in \mathbf{T}_\varphi$ for $\varphi \in \Phi$ and all of φ's variables appear in $\mathbf{var}_v(\mathbf{dom}_v)$, then A must satisfy φ.*

The instance of S-SAT that we construct, must have a list of satisfying assignments for each test. Note that the size of this list is bounded by the number of $[r_L, d]$-LDFs which is $|\mathcal{F}|^{O(1)}$, i.e. polynomial in n. Having defined the tests in Ψ and the satisfying assignments for each test, it now only remains to specify the variables that each test accesses, i.e. define for each node v, the mapping $\mathbf{var}_v : \mathbf{dom}_v \to V_\Psi$.

9.2 Variables

We begin by defining the variable mappings for the root nodes of the trees in the forest. Recall that for the root node $root_\varphi$ of each tree $\mathbf{T}_\varphi$, we set $\mathbf{dom}_{root_\varphi} = \mathcal{F}^{d_0}$. Let $\hat{V}_\Phi \supset V_\Phi$ be the variables representing the low-degree-extension (Definition 8.2, with parameters $m = |V_\Phi|$, $d_0 = \lceil \log\log n \rceil$, and $\mathcal{H} \subset \mathcal{F}$ such that[1] $|\mathcal{H}|^{d_0} = |V_\Phi|$) of V_Φ, i.e. $\hat{V}_\Phi$ is a set of $|\mathcal{F}|^{d_0}$ variables each representing a distinct point in $\mathcal{F}^{d_0}$. We define the mapping $\mathbf{var}_{root_\varphi}$ as follows,

Definition 9.3 ($\mathbf{var}_{root_\varphi}$) *The bijection $\mathbf{var}_{root_\varphi} : \mathbf{dom}_{root_\varphi} \to \hat{V}_\Phi$ maps the points of $\mathbf{dom}_{root_\varphi} = \mathcal{F}^{d_0}$ to $\hat{V}_\Phi$ in the following manner. Take $\mathcal{H} \overset{def}{=} \{0, .., h-1\} \subset \mathcal{F}$ such that $|\mathcal{H}|^{d_0} = h^{d_0} = |V_\Phi|$ (i.e. $|\mathcal{H}| = |V_\Phi|^{\frac{1}{d_0}} = n^{c/d_0} = n^{c/\log\log n}$ and since $|\mathcal{F}| = n^{c_1/\log\log n}$ we have $|\mathcal{H}|^{c_1/c} = |\mathcal{F}|$). We define $\mathbf{var}_{root_\varphi}$ to be a bijection* independent *of φ, taking the points of $\mathcal{H}^{d_0} \subset \mathcal{F}^{d_0}$ to V_Φ, and the remaining points $\mathcal{F}^{d_0} \setminus \mathcal{H}^{d_0}$ to $\hat{V}_\Phi \setminus V_\Phi$.*

Note that for every $\varphi \in \Phi$ the points of $\mathbf{dom}_{root_\varphi}$ were mapped to *the same* variables, hence each of the $|\mathcal{F}|^{d_0}$ variables in $\hat{V}_\Phi$ has $|\Phi|$ pre-images (so far).

[1] If $\sqrt[d_0]{m}$ is not an integer, we add dummy variables to V_Φ.

For simplicity we assume that for each $\varphi \in \Phi$, the points mapped to φ's variables are in general position (i.e. they span a $(D-1)$-dimensional affine-subspace of $\mathcal{F}^{d_0}$), otherwise, we choose an arbitrary $(D-1)$-dimensional affine subspace containing these points.

Before we continue to define the mappings $\mathbf{var}_v$ for non-root nodes, let us examine the purpose of these mappings. Picture a super-assignment to the tests of Ψ, as a labeling of each leaf in the forest by a super-LDF. We will prove (see Lemma 10.4) that such an assignment, if consistent and of low-norm, 'induces' a low-norm super-LDF for the domain of each *internal* node, and in particular – a low-norm super-LDF $\mathcal{G}$ for the 'root-domain', $\mathcal{F}^{d_0}$. We now use the fact that the variables representing this root-domain are common to the roots of all $\mathbf{T}_\varphi$'s, to interpret $\mathcal{G}$ as a global assignment for the variables in V_Φ. Namely, we will show that any LDF that appears in $\mathcal{G}$ with a non-zero coefficient assigns V_Φ values that satisfy most of the tests in Φ.

The idea behind the CR-forest is that the domain $\mathbf{dom}_u$ of a node u is 'represented' by its offspring' domains. u's domain's points are distributed among the domains of each of u's offspring. The aforementioned Lemma 10.4 will show how to join the LDFs of u's offspring into one LDF for u. The advantage we gain by representing one LDF over u's domain by many LDFs over u's offspring' domains is that we can enforce the degree of the LDFs in the leaves to be very low, compared to the degree of the LDF on the root that they represent (the dimension of the LDFs is maintained low as well). Therefore the list of satisfying assignments for the tests in Ψ (corresponding to LDFs on the leaves' domains) is not too long. We can afford to list all LDFs (i.e. satisfying assignments) only when the degree (and dimension) of the LDFs is small enough, because for a higher degree the length of the list would not be polynomial in n.

The key to understanding the construction is to see how a node u is 'represented' by its offspring. Pictorially, u's domain's points are distributed among the domains of u's offspring, each offspring v receives a slice of u's domain. Some of v's points correspond to v's slice of u's variables. The rest of v's points are some (low-degree) encoding or extension of these points.

Consider a non-root node v, and denote its parent by u. Assuming $\mathbf{var}_u$ is already defined, we now specify the mapping $\mathbf{var}_v : \mathbf{dom}_v \to V_\Psi$. Some (exactly $|\mathcal{F}|^{D+2}$) of the points in $\mathbf{dom}_v$ 'represent' points from $\mathbf{dom}_u$, and will thus be mapped to u's variables ($\mathbf{var}_u(\mathbf{dom}_u)$). The rest of the points in $\mathbf{dom}_v$ will be mapped to fresh new variables $V_v \subset V_\Psi$ ($|V_v| = |\mathcal{F}|^d - |\mathcal{F}|^{D+2}$) associated with the node v. Only points in domains of nodes in v's sub-tree may be mapped to V_v. For uniformity of notation, we define $V_{root_\varphi} \overset{def}{=} \hat{V}_\Phi$, for every root $root_\varphi$, again stressing the fact that the roots of all of the trees share the same variables. Altogether

$$V \overset{def}{=} V_\Psi = \bigcup_{\substack{v \in \mathbf{T}_\varphi \\ \varphi \in \Phi}} V_v \ .$$

u's variables are distributed among its offspring by letting each offspring v of u 'represent' an affine sub-space $\mathcal{C}_v \subset \mathbf{dom}_u$ of dimension $D+2$ (a $(D+2)$-cube). More formally, we label (as specified later in Section 9.3) each offspring v of u by a $(D+2)$-cube $\mathcal{C}_v \subset \mathbf{dom}_u$. We represent a cube $\mathcal{C}_v$ by $D+3$ points $x_0, .., x_{D+2}$ such that $\mathcal{C}_v = x_0 + \text{span}(x_1, .., x_{D+2})$ (this yields a natural way of viewing $\mathcal{C}_v$ as $\mathcal{F}^{D+2}$).

We embed all points of the cube $\mathcal{C}_v \subset \mathbf{dom}_u$ into the domain $\mathbf{dom}_v$ by the embedding extension mapping, defined above in Section 8.2, $E_{b_i} : \mathcal{C}_v \to \mathbf{dom}_v$ (the parameter b_i depends on the level $i \geq 1$ of the node v, and is specified shortly below). Via this mapping, we can transform LDFs on the cube $\mathcal{C}_v$ to *lower-degree* LDFs on the domain $\mathbf{dom}_v$. This will allow us to represent a satisfying assignment to Φ by $[r_i, d]$-LDFs on the domains $\mathbf{dom}_v$ of level-i nodes (the degree r_i will be defined below). The construction is aimed to *lower* the degree r_i of the LDFs, from $r_0 \overset{def}{=} |\mathcal{H}| \approx n^{1/\log\log n}$ to $r_L = O(1)$.

We think of the point $y = E_{b_i}(x) \in \mathbf{dom}_v$ as 'representing' the point $x \in \mathcal{C}_v \subset \mathbf{dom}_u$, and define $\mathbf{var}_v : \mathbf{dom}_v \to V_\Psi$ as follows,

Definition 9.4 ($\mathbf{var}_v$, for a non-root node v) *Let v be a non-root node, let u be v's parent, and let $\mathcal{C}_v \subset \mathbf{dom}_u$ be the label attached to v (the* label *of a node is defined below, Definitions 9.5,9.6). For each point $y \in E_{b_i}(\mathcal{C}_v) \subset \mathbf{dom}_v$ define $\mathbf{var}_v(y) \overset{def}{=} \mathbf{var}_u(E_{b_i}^{-1}(y))$, i.e. points that 'originated' from $\mathcal{C}_v$ are mapped to the previous-level variables, that their pre-images in $\mathcal{C}_v$ were mapped to. For each 'new' point $y \in \mathbf{dom}_v \setminus E_{b_i}(\mathcal{C}_v)$ we define $\mathbf{var}_v(y)$ to be a distinct variable from V_v.*

The parameters used for the embedding extension mappings E_{b_i} are $t = D+2$, $k = d/t$. We set $r_0 = |\mathcal{H}| = |\mathcal{F}|^{c/c_1}$ and r_{i+1} and b_{i+1} ($i \geq 0$) are defined by the following recursive formulas:

$$\begin{aligned} b_{i+1} &= \left\lceil \sqrt[4]{r_i(D+2)+1} \right\rceil \\ r_{i+1} &= b_{i+1} - 1 \end{aligned}$$

(we will show in Section 10.1 that b_i, r_i decrease until for some $L < \log\log n$, $r_L \leq 2(D+2) = O(1)$).

In order to complete the description of the test-system, we now only need to describe the cube-labeling of all of the offspring of each node. This will describe how the representation of a node u is distributed among its offspring.

9.3 Labeling Nodes

We define the *offspring-labels* of a node u, thereby completing the description of the construction. As described above, each offspring of the node u 'represents' an affine subspace in the domain $\mathbf{dom}_u$, i.e. the variables of u's offspring represent an encoding of u's variables.

This representation has some error. To control this error, we proportion the offspring so that more important variables are represented by more offspring. Roughly speaking, the 'importance' of a variable $x \in V$ is determined by how high up (towards the root) in the tree this variable appears. The closer the variable is to the root, the more information it represents about Φ's original variables, V_Φ.

Let us begin by defining the labels of the offspring of a root node $root_\varphi$. The tests at the leaves of the tree $\mathbf{T}_\varphi$ represent the test $\varphi \in \Phi$. Therefore, the variables that φ depends on are 'very important' to represent. We define the offspring-labels of $root_\varphi$ so that every offspring represents these variables,

Definition 9.5 (offspring-labels for a root node) *Let*

$$\mathbf{dst}(root_\varphi) \overset{def}{=} \left\{ x \in \mathbf{dom}_{root_\varphi} \,\middle|\, \varphi \textit{ depends on } \mathbf{var}_{root_\varphi}(x) \right\}$$

be a set of distinguished points for $root_\varphi$ (recall our assumption that $\mathbf{dst}(root_\varphi)$ is a set of exactly D points in general position). We label each offspring of $root_\varphi$ by a distinct cube from the following set:

$$\mathbf{labels}(root_\varphi) \overset{def}{=} \left\{ \mathcal{C} \textit{ is a } (D+2)\textit{-cube in } \mathcal{F}^{d_0} \,\middle|\, \mathcal{C} \supset \mathbf{dst}(root_\varphi) \right\}$$

The number of labels $|\mathbf{labels}(root_\varphi)| = B(\mathcal{F}^{d_0}, D+2, D-1)$ is the number of $(D+2)$-cubes containing the $(D-1)$-cube spanned by the points mapped to φ's variables (assuming, as mentioned above, that these points are in general position).

For a general non-root node $v \in \mathbf{T}_\varphi$, we use a mechanism of 'distinguished-points' to promote the importance of certain points more than others. For each offspring v of u we define, hand in hand with v's label, a set of distinguished points $\mathbf{dst}(v)$. v will 'represent' these points in the sense that the descendants of v will represent $\mathbf{dst}(v)$ with special care. For a general non-root node $v \in \mathbf{T}_\varphi$, we consider two levels of 'important' variables: (1) variables that belong to some ancestor (direct and indirect) of v (there are $|\mathcal{F}|^{D+2}$ such variables, all mapped from v's parent) and (2) variables mapped from the distinguished points of v (there will always be exactly one or D such variables). The node v will correspondingly have two equal-weight sets of offspring,

Definition 9.6 (offspring-labels for non-root nodes) *Let v be a non-root node. We define two multi-sets of offspring-labels for v. For each variable $\mathtt{x} \in \mathbf{var}_v(\mathbf{dom}_v) \setminus V_v$, i.e. $\mathtt{x}$ that belongs to some ancestor of v, we define*

$$\mathbf{labels}_{\mathtt{x}}(v) \overset{def}{=} \{\mathcal{C} \subset \mathbf{dom}_v \textit{ is a } (D+2)\textit{-cube} \,|\, \mathtt{x} \in \mathbf{var}_v(\mathcal{C})\}$$

we then take $\mathbf{labels}_1(v)$ to be the multi-set

$$\mathbf{labels}_1(v) \overset{def}{=} \bigcup_{\mathtt{x} \in \mathbf{var}_v(\mathbf{dom}_v) \setminus V_v} \mathbf{labels}_{\mathtt{x}}(v)$$

For every offspring w of v, labelled by a cube from $\mathbf{labels}_{\mathtt{x}}(v)$, we define $\mathbf{dst}(w)$ to be the singleton set consisting of the point in $\mathbf{dom}_w$ that is mapped to x*, i.e.* $\mathbf{dst}(w) \overset{def}{=} \{\mathbf{var}_w{}^{-1}(\mathtt{x})\} \subset \mathbf{dom}_w$.

The second multi-set (actually set) of offspring-labels is devoted to representing the distinguished points of v. We simply take

$$\mathbf{labels}_2(v) \overset{def}{=} \{\mathcal{C} \subset \mathbf{dom}_v \textit{ is a } (D+2)\textit{-cube} \,|\, \mathcal{C} \supset \mathbf{dst}(v)\}$$

For each offspring w of v labelled by a cube from $\mathbf{labels}_2(v)$, we set $\mathbf{dst}(w) \overset{def}{=} E_{b_i}(\mathbf{dst}(v))$ (where i is w's level in the tree), i.e. w distinguishes the same set of variables as v.

The final multi-set $\mathbf{labels}(v)$ is the union of $\mathbf{labels}_1(v)$ and $\lfloor M \rfloor$ copies of $\mathbf{labels}_2(v)$, where the number $M = |\mathcal{F}|^{D+2} B(\mathcal{F}^d, D+2, 0)/B(\mathcal{F}^d, D+2, |\mathbf{dst}(v)|-1)$ is chosen so that at least half of the labels are from $\mathbf{labels}_1$, and at least a third of the labels are from $\mathbf{labels}_2$.

9.4 Construction Size

Recall that we defined $d_0 \overset{def}{=} \lceil \log\log n \rceil$ and $d \overset{def}{=} 4D+8$. We also set $r_0 = |\mathcal{H}| = |\mathcal{F}|^{c/c_1} = n^{c/\log\log n}$, and defined $b_{i+1} = \left\lceil \sqrt[4]{(D+2)r_i + 1} \right\rceil$ and $r_{i+1} = b_{i+1} - 1$ for every $i \geq 0$.

We claim that indeed $r_L = O(1)$ for some $L \leq \log\log n$. For this purpose we prove by simple induction that $r_i \leq \max(\left\lceil r_0^{1/2^i} \right\rceil, 2(D+2))$. For r_0 this indeed holds, and assuming it true for r_i we have that if $r_i > 2(D+2)$ and $b_{i+1} > 2$, then

$$r_{i+1} < b_{i+1} = \left\lceil \sqrt[4]{(D+2)r_i + 1} \right\rceil < \left\lceil \sqrt[4]{2r_i(D+2)} \right\rceil \leq \left\lceil (r_i)^{2/4} \right\rceil \leq \lceil \sqrt{r_i} \rceil \leq \left\lceil (r_0)^{1/2^{i+1}} \right\rceil .$$

We set L to be the first index for which $r_L \leq 2(D+2) = O(1)$. Obviously, until that point r_i, b_i decrease monotonically, and since $r_0 = 2^{c \log n / \log\log n}$, $L \leq \lfloor \log(c \log n / \log\log n) \rfloor + 1 < \log\log n$. This also implies that $b_i > 2$ for all $0 \leq i \leq L$, and completes the induction.

The Range of the Tests. The tests of the test-system range over $[r_L, d]$-LDFs. The number of monomials of degree $r_L \leq 2(D+2) = O(1)$, and dimension $d = 4D+8 = O(1)$ is bounded by $(r_L+1)^d = O(1)$. The number of $[r_L, d]$-LDFs is hence bounded by $|\mathcal{F}|^{O(1)} < O(n)$ and therefore the range of the tests is polynomial in n.

The Number of Tests and Variables. It is only left to verify that the size of the forest is polynomial. We have $|\Phi| = n$ trees, so let's verify that the number of nodes in each tree is polynomially-bounded.

Consider a tree $\mathbf{T} = \mathbf{T}_\varphi \in \mathbf{F}_n(\Phi)$. $root_\varphi$ has $B(\mathcal{F}^{d_0}, D+2, D-1) \leq |\mathcal{F}|^{3d_0} = n^{O(1)}$ offspring and each node in level i $(0 < i < L)$ has $2\,|\mathcal{F}|^{D+2} \cdot B(\mathcal{F}^d, D+2, 0) = |\mathcal{F}|^{O(1)}$ offspring. Altogether the number of nodes in $\mathbf{T}$ is bounded by

$$n^{O(1)} \cdot \prod_{i=1}^{L} |\mathcal{F}|^{O(1)} = n^{O(1)} \cdot |\mathcal{F}|^{O(L)} = n^{O(1)} \cdot (2^{\log n / \log\log n})^{O(\log\log n)} = n^{O(1)}$$

Hence the number of tests in Ψ is polynomial, and the number of variables is $\leq |\mathcal{F}|^d \cdot |\Psi| = n^{O(1)}$.

Chapter 10

Correctness of the Reduction

In this chapter we prove the completeness and soundness of the reduction presented in the previous chapter.

10.1 Completeness

Lemma 10.1 (Completeness) *If there is an assignment $\mathcal{A} : V_\Phi \to \{\texttt{true}, \texttt{false}\}$ satisfying all of the tests in Φ, then there is a natural assignment $\mathcal{A}_\Psi : V_\Psi \to \mathcal{F}$ satisfying all of the tests in Ψ.*

Of course, this assignment $\mathcal{A}_\Psi$ is equivalent to a consistent natural super-assignment. We extend $\mathcal{A}$ following the rationale of the construction, by taking its low-degree-extension to the variables $\hat{V}_\Phi$, and then repeatedly taking the embedding extension of the previous-level variables, until we've assigned all of the variables in the system. More formally,

Proof: We construct an assignment $\mathcal{A}_\Psi : V_\Psi \to \mathcal{F}$. We first set (for every $\varphi \in \Phi$) P_{root_φ} to be the low degree extension (see Definition 8.2) of $\mathcal{A}$ (we think of $\mathcal{A}$ as assigning each variable a value in $\{0,1\} \subset \mathcal{F}$ rather than $\{\texttt{true}, \texttt{false}\}$, see discussion in the beginning of Chapter 9). We proceed to inductively obtain $[r_i, d]$-LDFs $P_v : \mathbf{dom}_v \to \mathcal{F}$ for every level-i ($i > 0$) node v of every tree in the CR-forest, as follows. Assume we've defined an $[r_i, d]$-LDF (an $[r_i, d_0]$-LDF in case $i = 0$) P_u consistently for all level-i nodes, and let v be an offspring of u, labelled by $\mathcal{C}_v$. The restriction $f = P_u|_{\mathcal{C}_v}$ of P_u to the cube $\mathcal{C}_v$ is a $[r_i(D+2), D+2]$-LDF. f can be written as a $[\left\lceil \sqrt[4]{r_i(D+2)+1} \right\rceil - 1,\ 4D+8]$-LDF f_{ext} over the larger domain $\mathcal{F}^d$, as promised by Proposition 8.3 taking $k = 4$. We define $P_v = f_{\mathbf{ext}}$ to be that $[r_{i+1}, d]$-LDF (recall that $d = 4D+8$ and $r_{i+1} = b_{i+1} - 1 = \left\lceil \sqrt[4]{r_i(D+2)+1} \right\rceil - 1$).

Finally, for a variable $\mathtt{x} \in \mathbf{var}_v$, $\mathtt{x} = \mathbf{var}_v(x)$, we set $\mathcal{A}_\Psi(\mathtt{x}) \overset{def}{=} P_v(x)$. The construction implies that there are no collisions, i.e. $\mathtt{x}' = \mathbf{var}_{v'}(x') = \mathbf{var}_v(x) = \mathtt{x}$ implies $P_v(x) = P_{v'}(x')$. ■

10.2 Soundness

In this section we show that a 'no' instance of PCP is always mapped to a 'no' instance of S-SAT. We assume that the constructed S-SAT instance has a consistent super-assignment of norm $\leq g$, and show that Φ – the PCP test system we started with – is satisfiable.

Lemma 10.2 (Soundness) *Let $g \overset{def}{=} |\mathcal{F}|^{c_g}$ where $c_g > 0$ is some small enough constant, say $c_g = 1/1000$. If there exists a non-trivial consistent super-assignment of norm $\leq g$ for Ψ, then Φ is satisfiable.*

Let us first sketch a brief outline of the proof. The proof follows the structure of the trees underlying the construction. Since the tree structure is different for the first level nodes and for all other levels, we divide the proof accordingly.

We begin with a few definitions preparing for the proof itself. We then state Lemma 10.4 that encapsulates the inductive part, handling all internal nodes in levels ≥ 1 of the tree, and proving that a non-trivial consistent super-assignment at the leaves can be decoded into "consistent" super-LDFs on "most" internal nodes. Relying on this lemma, we proceed to prove the soundness lemma (Lemma 10.2). The heart of the proof is a consistency lemma (Lemma 10.5) that allows us to combine "consistent" super-LDFs on domains of offspring of a node into one super-LDF on the domain of that node. We use this lemma to combine the super-LDFs on the root's offspring (i.e. level-1 nodes) into one global super-LDF on the common domain $\mathcal{F}^{d_0}$, and from it deduce an assignment satisfying the original PCP test-system Φ.

We then return to the inductive proof of Lemma 10.4 again relying on the same consistency lemma (Lemma 10.5) for the inductive step.

The proof of the consistency lemma (Lemma 10.5) itself follows in Section 10.3.

Proof: Let $\mathcal{SA}$ be a non-trivial consistent super-assignment for Ψ, of norm $\|\mathcal{SA}\| \leq g$. It induces (by projection) a super-assignment to the variables

$$m : V_\Psi \longrightarrow \mathbb{Z}^{|\mathcal{F}|}$$

i.e. for every variable $\mathtt{x} \in V_\Psi$, m assigns a vector $\pi_{\mathtt{x}}(\mathcal{SA}(\psi))$ of integer coefficients, one per value in $\mathcal{F}$ where ψ is some test depending on $\mathtt{x}$. Since $\mathcal{SA}$ is consistent, m is well defined (independent of the choice of test ψ). Alternatively, we view m as a labeling of the points $\bigcup_{v \in \mathbf{F}_n(\Phi)} \mathbf{dom}_v$ by a 'super-value' – a formal linear combination of values from $\mathcal{F}$. The label of the point $x \in \mathbf{dom}_v$ for some $v \in \mathbf{F}_n(\Phi)$, is simply $m(\mathbf{var}_v(x))$, and with a slight abuse of notation, is sometimes denoted $m(x)$. m is used as the "underlying point super-assignment" for the rest of the proof, and will serve as an anchor by which we test consistency. Since $\mathcal{SA}$ is non-trivial, $m(x) \neq \vec{0}$ for every x.

For a node u, we denote by $\mathrm{Avg}(u)$ the average of $\|\mathcal{SA}(\psi_v)\|$ over the leaves v in u's sub-tree. We will show that whenever $\mathrm{Avg}(u)$ is not too high for a node u, then u's subtree is, in a sense, consistent. We thus define a 'good' node as one having a low average norm on its subtree's leaves:

Definition 10.1 (Good Nodes) *Fix $C > 0$ large enough, e.g. $C = 301$. A node u in level i in the CR-forest is said to be* good *if*

$$\mathrm{Avg}(u) \le g_i \overset{def}{=} g \cdot C^{i+1}.$$

We denote by $\mathbf{nodes}_i^*$ the set of good nodes in level i.

For any node v, denote by $\mathbf{ofsp}(v)$ the set of v's direct offspring. It is easy to see that most offspring of a good node are themselves good:

Proposition 10.3 *If $u \in \mathbf{nodes}_i^*$, then*

$$\Pr_{v \in \mathbf{ofsp}(u)}(v \in \mathbf{nodes}_{i+1}^*) \ge 1 - 1/C$$

Proof: u is good, thus by definition $\mathrm{Avg}(u) \le g_i$. Had u more than $1/C$ bad offspring, then the total average of $\mathcal{SA}(\psi_v)$ on its sub-tree would be $> g_{i+1} \cdot \frac{1}{C} = g_i$. ■

The central task of our proof is to show that a consistent low-norm super-assignment to the tests at the leaves induces a low-norm super-LDF on the root domain. The key step in this proof is the inductive step showing that if, for a node v, almost all of its offspring have a low-norm super-LDF that is consistent with m, then we can deduce such a super-LDF $\mathcal{G}_v$ over $\mathbf{dom}_v$.

It turns out that for a general node u we cannot always deduce a super-LDF agreeing with m on every point in $\mathbf{dom}_u$ (a counter-example can be constructed). Instead, for good nodes u we show that there exists a super-LDF $\mathcal{G}_u$ over $\mathbf{dom}_u$ that agrees with *almost* all of the super-LDFs on u's offspring. By 'agrees' we mean that (suppose v is an offspring of u labelled by $\mathcal{C}_v$) the parent super-LDF $\mathcal{G}_u$ projected on the points of $\mathcal{C}_v$ equals $\mathcal{G}_v$ projected on the manifold points $E_{b_i}(\mathcal{C}_v)$.

The consistency with m will then follow inductively from the fact that the offspring' super-LDFs were consistent with m. The importance of consistency with m is not the same for all points. For certain points (e.g. those mapped to variables from V_Φ) we cannot allow any inconsistency, while for others we can allow some small error. For every node v we consider, as mentioned before in the construction, two types of special points: The distinguished points $\mathbf{dst}(v)$, and the manifold points,

Definition 10.2 (Manifold Points) *For a non-root node v labelled by $\mathcal{C}_v$, we define the manifold points $\mathbf{manf}(v) \overset{def}{=} E_{b_i}(\mathcal{C}_v)$, where i is v's level in the tree.*

These are the points that originate from all of v's ancestors.

We rely on the manner in which the offspring were proportioned to deduce a high level of consistency for these special points. We can now state the key inductive lemma (Lemma 10.4) showing how consistency follows from having a low-norm super-assignment. This lemma relies heavily on the precise structure of the forest, and shows that for every good node u, there is a super-LDF on u's domain that is 'almost' consistent with "the anchor" m. The lemma is proved inductively, constructing u's super-LDF from the super-LDFs of u's good offspring. We will later want to construct from these super-LDFs an assignment that satisfies more than half of the tests in Φ. For this purpose, we need the super-LDFs along the way to be legal,

Definition 10.3 (Legal) *An LDF P is called* legal for a node $v \in \mathbf{T}_\varphi$ *(for some $\varphi \in \Phi$), if it satisfies φ in the sense that if φ's variables have pre-images $x_1, .., x_D \in \mathbf{dom}_v$, then $P(x_1), .., P(x_D)$ satisfy φ. A super-LDF $\mathcal{G}$ is called* legal for $v \in \mathbf{T}_\varphi$ *if for every LDF P appearing in $\mathcal{G}$, P is legal for $v \in \mathbf{T}_\varphi$.*

Lemma 10.4 *Let $u \in \mathbf{nodes}_i^*$ for some $i \geq 1$, and set $\alpha = 3/C$ and $A = 4 \cdot (D+2)^3 = O(1)$. There exists a legal super-LDF $\mathcal{G}_u$ of degree at most $\tilde{r}_i \overset{def}{=} A^{L-i} \cdot (r_i + 1)$ and of norm $\leq 2^{L-i} \cdot \mathrm{Avg}(u)$ that agrees with m on $\mathbf{dst}(u)$ and on $1-\alpha$ fraction of the points in $\mathbf{manf}(u)$, i.e.*

$$\forall x \in \mathbf{dst}(u) \quad \pi_x(\mathcal{G}_u) = m(x)$$

and

$$\Pr_{x \in \mathbf{manf}(u)} (\pi_x(\mathcal{G}_u) = m(x)) \geq 1 - \alpha$$

This lemma is the key to our construction. Its proof shows how consistent LDFs on offspring nodes induce an LDF on the parent. Before we prove this lemma, which is somewhat technical, let us first use it to complete the proof of Lemma 10.2.

Applying Lemma 10.4 for level-1 nodes, we deduce that $\mathcal{SA}$ induces a legal super-LDF $\mathcal{G}_v$ of degree $\tilde{r}_1$ and with $\|\mathcal{G}_v\| \leq 2^{L-1}\mathrm{Avg}(v)$, for every node v in $\mathbf{nodes}_1^*$. We now join these super-LDFs into one legal super-LDF over the root domain $\mathcal{F}^{d_0}$, and then deduce a satisfying assignment to the tests in Φ from this super-LDF. Let $v \in \mathbf{nodes}_1^*$ be an offspring of $root_\varphi$ for some $\varphi \in \Phi$, and let $\mathcal{C}_v \subset \mathbf{dom}_{root_\varphi}$ be the cube labeling v. We would like to view $\mathcal{G}_v$ as a super-LDF over $\mathcal{C}_v$, by restricting the LDFs in $\mathcal{G}_v$ to the manifold $\mathbf{manf}(v) \subset \mathbf{dom}_v$ that represents $\mathcal{C}_v$. For every $[\tilde{r}_1, d]$-LDF $P : \mathbf{dom}_v \to \mathcal{F}$, define $\tilde{P}$ as the $[d\tilde{r}_1 \cdot (b_1)^3, D+2]$-LDF which is defined as P's restriction to the manifold $\mathbf{manf}(v)$:

$$\forall x \in \mathcal{C}_v \subset \mathbf{dom}_{root_\varphi} : \quad \tilde{P}(x) \overset{def}{=} P(E_{b_1}(x))$$

(Note that since P's total degree is $\leq \tilde{r}_1 \cdot d$, the total degree of $\tilde{P}$ is $\leq \tilde{r}_1 \cdot d \cdot (b_1)^3$ because the degree of E_{b_1} is $(b_1)^{4-1} = (b_1)^3$). For every LDF $P : \mathbf{dom}_v \to \mathcal{F}$, $\mathcal{G}_v$ assigns an integer value $\mathcal{G}_v[P]$. We define the super-LDF $\tilde{\mathcal{G}}_v$ to be the same formal linear combination as $\mathcal{G}_v$, replacing each LDF P with $\tilde{P}$:

$$\tilde{\mathcal{G}}_v[\tilde{Q}] = \sum_{P,\ \tilde{P}=\tilde{Q}} \mathcal{G}_v[P]$$

In other words, the super-LDF $\tilde{\mathcal{G}}_v$ is simply the restriction (and re-parameterization) of $\mathcal{G}_v$ to the manifold $\mathbf{manf}(v)$, as discussed in Section 8.2. The total degree of $\tilde{\mathcal{G}}_v$ is $\tilde{r}_1 d \cdot (b_1)^3$.

Let $root_\varphi$ be a good root node. Since the average norm of all the tests is $\leq g$, and a good root node is by definition one with $\mathrm{Avg}(root_\varphi) \leq g \cdot C$, there are at least $1 - 1/C = 1 - \alpha/3$ such nodes. For every good offspring v of $root_\varphi$, Lemma 10.4 guarantees that $\Pr_{x \in \mathcal{C}_v}\left(\pi_x(\tilde{\mathcal{G}}_v) = m(x)\right) \geq 1 - \alpha$, and that for every $x \in \mathbf{dst}(v)$, $\pi_x(\tilde{\mathcal{G}}_v) = m(x)$. Given this assignment of super-LDF $\tilde{\mathcal{G}}_v$ per label $\mathcal{C}_v \subset \mathbf{dom}_{root_\varphi} = \mathcal{F}^{d_0}$, we would like to use the fact that these super-LDFs are consistent with m to deduce the existence of some global super-LDF on $\mathcal{F}^{d_0}$ (that is also consistent with m). The following consistency lemma, when applied for $u = root_\varphi$, will imply just that.

Lemma 10.5 (Consistency Lemma) *Let $u \in \mathbf{nodes}^*_i$ for some $0 \leq i < L$. Define $\mathcal{S}^*$ to be the multi-set of cubes that label good offspring of u, $\mathcal{S}^* \overset{def}{=} \left\{\mathcal{C}_v \in \mathbf{labels}(u) \mid v \in \mathbf{nodes}^*_{i+1}\right\}$. If for every good offspring v of u there is a super-LDF $\tilde{\mathcal{G}}_v$ over $\mathcal{C}_v$, of total degree $\leq r = \tilde{r}_i/(D+2)$ and norm $\|\tilde{\mathcal{G}}_v\| \leq 2^{L-i-1} \cdot \mathrm{Avg}(v)$, such that*

$$\Pr_{x \in \mathcal{C}_v}\left(\pi_x(\tilde{\mathcal{G}}_v) = m(x)\right) \geq 1 - \alpha$$

then there is a super-LDF $\mathcal{G}_u$ over $\mathbf{dom}_u$ of total degree $\tilde{r}_i = r(D+2)$ and norm $\|\mathcal{G}_u\| \leq 2^{L-i} \cdot \mathrm{Avg}(u)$ that obeys

$$\Pr_{\mathcal{C}_v \in \mathcal{S}^*}\left(\pi_{\mathcal{C}_v}(\mathcal{G}_u) = \tilde{\mathcal{G}}_v\right) \geq 1 - \alpha/6$$

We defer the proof of this lemma to the next section, and continue with the proof of Lemma 10.2. As previously mentioned, the super-LDFs $\tilde{\mathcal{G}}_v$ obtained for the $\mathcal{C}_v$s were of total degree $\tilde{r}_1 \cdot d(b_1)^3$, hence the degree is:

$$\begin{aligned}
\tilde{r}_1 \cdot d(b_1)^3 &= A^{L-1}(r_1+1) \cdot d(b_1)^3 = A^{L-1} \cdot d \cdot (b_1)^4 = \\
&= A^{L-1} \cdot 4(D+2) \cdot (r_0(D+2)+1) \\
&< A^{L-1} \cdot 4(D+2)^2(r_0+1) = A^L(r_0+1)/(D+2) = \tilde{r}_0/(D+2)
\end{aligned}$$

using the definitions $A = 4(D+2)^3$, $\tilde{r}_i = A^{L-i}(r_i+1)$, $b_{i+1} = \sqrt[4]{r_i(D+2)+1}$ and $r_{i+1} = b_{i+1} - 1$. Hence we obtain from the consistency lemma a global super-LDF $\mathcal{G}_\varphi$ of degree $\tilde{r}_0$ over $\mathcal{F}^{d_0}$ that agrees with $\mathcal{G}_v$ for $1 - \alpha/6$ of the good offspring v of $u = root_\varphi$.

We next show that $\mathcal{G}_\varphi = \mathcal{G}_{\varphi'}$ for every $\varphi \neq \varphi'$ whose corresponding nodes $root_\varphi$ and $root_{\varphi'}$ are both good. Choose a random cube $\mathcal{C}_v \in \mathbf{labels}(root_\varphi)$ and a random point

$x \in \mathcal{C}_v \subset \mathcal{F}^{d_0}$. We claim that $\Pr_{x,\mathcal{C}_v}(\pi_x(\tilde{\mathcal{G}}_v) = m(x)) \geq 1 - 2\alpha$. By Proposition 10.3 the probability that $\mathcal{C}_v$ is not good is $\leq \alpha/3$. If $\mathcal{C}_v$ is good, the above Lemma 10.5 tells us that with probability at most $\alpha/6$, $\pi_{\mathcal{C}_v}(\mathcal{G}_\varphi) \neq \tilde{\mathcal{G}}_v$ (altogether we have that with probability $\geq 1-\alpha/3-\alpha/6 \geq 1-\alpha$ over the cubes in $\mathbf{labels}(\varphi)$, $\pi_{\mathcal{C}_v}(\mathcal{G}_\varphi) = \tilde{\mathcal{G}}_v$). Now, by Lemma 10.4, for any good $\mathcal{C}_v$, $\Pr_{x\in\mathcal{C}_v}(\pi_x(\tilde{\mathcal{G}}_v) \neq m(x)) \leq \alpha$. For all otherwise chosen points, we have $\pi_x(\mathcal{G}_\varphi) = m(x)$, and the claim is proven.

These points constitute roughly[1] a $1-2\alpha$ fraction of $\mathcal{F}^{d_0}$. Hence $\mathcal{G}_\varphi$ and $\mathcal{G}_{\varphi'}$ agree with m on the same $\geq 1 - 4\alpha > 1/2$ fraction of the points. Thus, the super-LDF $\mathcal{G}_\varphi - \mathcal{G}_{\varphi'}$ (subtraction is defined as subtraction of the coefficient vectors) is zero when projected on more than half of the points.

Now, utilizing the fact that $\|\mathcal{G}_\varphi - \mathcal{G}_{\varphi'}\| \leq \|\mathcal{G}_\varphi\| + \|\mathcal{G}_{\varphi'}\| \leq 2^{L+1}g$, and by the low-ambiguity property (see Proposition 8.1) the fraction of ambiguous points (the only candidates on which the projection can be zero) is bounded by

$$\mathrm{amb}(\tilde{r}_0, d_0, 2^{L+1}g) < 2^{2(L+1)}g^2\frac{\tilde{r}_0 d_0}{|\mathcal{F}|} \leq 2^{2(L+1)} \cdot A^L \cdot d_0 \cdot |\mathcal{F}|^{2c_g+c/c_1-1} \ll 1/2$$

Thus, we deduce that $\mathcal{G}_\varphi = \mathcal{G}_{\varphi'}$. In addition, $\mathcal{G}_\varphi$ must be non-trivial since $m(x) \neq \vec{0}$ for every x.

We choose an arbitrary LDF P that appears in $\mathcal{G} \overset{def}{=} \mathcal{G}_{\varphi'} \neq \vec{0}$ for some good $root_{\varphi'}$, and define an assignment $\mathcal{A}_P : V_\Phi \to \{\texttt{true}, \texttt{false}\}$ for the variables of Φ as follows. For each $v \in V_\Phi$, we define $\mathcal{A}_P(v) \overset{def}{=} \texttt{true}$ iff $P(x) = 0$ on the corresponding point $x = {\mathbf{var}_{root_\varphi}}^{-1}(v)$ (see Definition 9.3), and $\mathcal{A}_P(v) \overset{def}{=} \texttt{false}$ otherwise.

The fraction of tests $\varphi \in \Phi$ for which $root_\varphi$ is good is at least $1 - 1/C > 1/2$ (because the total average of $\mathrm{Avg}(root_\varphi)$ over all $\varphi \in \Phi$ is g, and a good root node $root_\varphi$ is defined as a node with $\mathrm{Avg}(root_\varphi) \leq C \cdot g$).

We will show that $\mathcal{A}_P$ satisfies φ for every good node $root_\varphi$, and thus Φ is totally satisfiable. Let $\varphi \in \Phi$ be such that $root_\varphi$ is a good node. By the above consistency lemma, we know that for $1 - \alpha/6$ of the good offspring v of $root_\varphi$, $\pi_{\mathcal{C}_v}(\mathcal{G}) = \tilde{\mathcal{G}}_v$, and unless P is cancelled on $\mathcal{C}_v$, $P|_{\mathcal{C}_v}$ appears in $\tilde{\mathcal{G}}_v$. P will be cancelled on $\mathcal{C}_v$ only if there is another LDF Q appearing in $\mathcal{G}$ whose restriction to $\mathcal{C}_v$ equals P's restriction. For each Q the probability for this is bounded by (see Proposition 8.2) $\frac{\tilde{r}_0 d_0}{|\mathcal{F}|}$. Since there are no more than $2^K g$ possible LDFs Q that appear in $\mathcal{G}$, P is cancelled with probability $\leq \frac{\tilde{r}_0 d_0}{|\mathcal{F}|} \cdot 2^K g \ll 1/2$.

Thus there exists at least one good offspring v of $root_\varphi$ for which $\pi_{\mathcal{C}_v}(\mathcal{G}) = \tilde{\mathcal{G}}_v$ and $P|_{\mathcal{C}_v}$ appears in $\tilde{\mathcal{G}}_v$. Recall that the distinguished points of each offspring v of a root node $root_\varphi$, $\mathbf{dst}(v)$, are mapped to φ's variables. In addition, since v is good, Lemma 10.4 ensures that $\tilde{\mathcal{G}}_v$ is legal, i.e. for every Q appearing in $\tilde{\mathcal{G}}_v$, Q's restriction to φ's variables satisfies φ. It follows that φ is satisfied by $\mathcal{A}_P$. ■

[1] This procedure is *almost* equivalent to choosing a point uniformly at random, however there is a small (negligible) bias in favor of points in the span of $\mathbf{dst}(root_\varphi)$

This completes the proof of soundness, Lemma 10.2.

We now fill in the proof of Lemma 10.4.

Proof: (of Lemma 10.4) We prove this statement by induction on $L-i$. We ascend from the leaves to the top level, obtaining a super-LDF for each good node from the super-LDFs of its good offspring.

For obtaining the base of the induction ($i = L$), recall that for every leaf $u \in \mathbf{nodes}^*_L$, the test ψ_u is assigned a super-LDF $\mathcal{SA}(\psi_u)$. The definition of $\mathbf{nodes}^*_L$ implies $\|\mathcal{SA}(\psi_u)\| \leq g_L$. Since $\mathcal{SA}$ is a consistent super-assignment, $\mathcal{SA}(\psi_u)$ agrees with m on all of $\mathbf{dom}_u$ (in particular with all of $\mathbf{manf}(u)$ and $\mathbf{dst}(u)$), and thus the base of the induction is established.

To see the inductive step ($1 \leq i < L$), let $u \in \mathbf{nodes}^*_i$ be a good level-i node. By the inductive hypothesis for $L-i-1$, every good offspring v of u has a legal super-LDF $\mathcal{G}_v$ (of degree $\tilde{r}_{i+1}$) with norm $\leq 2^{L-i-1} \cdot \mathrm{Avg}(v)$ such that $\mathcal{G}_v$ agrees with m on $\mathbf{dst}(v)$ and on $1-\alpha$ of the points in $\mathbf{manf}(v)$.

Let v be a good offspring of u. We define $\tilde{\mathcal{G}}_v$, as in the proof of Lemma 10.2, to be the same linear combination as $\mathcal{G}_v$, taking LDFs $\tilde{P}$ instead of P, where $\tilde{P} : \mathcal{C}_v \to \mathcal{F}$ is defined by $\forall x \in \mathcal{C}_v \; \tilde{P}(x) \overset{def}{=} P(E_{b_{i+1}}(x))$. It follows from definition 8.3 and from the inductive hypothesis that $\tilde{\mathcal{G}}_v$ is a super-LDF of total degree $d\tilde{r}_{i+1} \cdot (b_{i+1})^3$ as before. As before, this is bounded by $\leq \tilde{r}_i/(D+2)$.

For any good label $\mathcal{C}_v$, let $x \in \mathcal{C}_v$, and let $y = E_{b_{i+1}}(x) \in \mathbf{manf}(v) \subset \mathbf{dom}_v$. Recall that we abbreviated $m(x)$ to mean $m(\mathbf{var}_v(x))$. Furthermore, we defined $\mathbf{var}_v(y) = \mathbf{var}_v(E_{b_{i+1}}(x)) = \mathbf{var}_u(x)$, thus $m(y) = m(x)$. Note also that by definition of $\tilde{\mathcal{G}}_v$, $\pi_x(\tilde{\mathcal{G}}_v) = \pi_y(\mathcal{G}_v)$.

By the inductive hypothesis, and since $E_{b_{i+1}}$ bijects $\mathcal{C}_v$ to $\mathbf{manf}(v)$, we have that the following equality holds both (1) for $1-\alpha$ of the points $x \in \mathcal{C}_v$, and (2) for every point $x \in \mathcal{C}_v$ such that $y = E_{b_{i+1}}(x) \in \mathbf{dst}(v)$:

$$\pi_x(\tilde{\mathcal{G}}_v) = \pi_y(\mathcal{G}_v) = m(y) = m(x)$$

By (1), and applying the consistency lemma (Lemma 10.5), we deduce a global super-LDF $\mathcal{G}_u$ over $\mathbf{dom}_u$ of norm $\|\mathcal{G}_u\| \leq 2^{L-i} \cdot \mathrm{Avg}(u)$ and of degree $\tilde{r}_i$ such that for $1-\alpha/6$ of the cubes in $\mathcal{S}^*$,

$$\pi_{\mathcal{C}_v}(\mathcal{G}_u) = \tilde{\mathcal{G}}_v \qquad (*)$$

This constitutes at least $1-\alpha/6-1/C = 1-\alpha/2$ of all the cubes in $\mathbf{labels}(u)$ (there are no more than $1/C$ no-good cubes, see Proposition 10.3). Now recall from the construction of the offspring-labels (see Definition 9.6) that at least one half of the offspring v of u are labelled by cubes in $\mathbf{labels}_1(u)$. We deduce that $1 - 2\cdot\alpha/2 = 1-\alpha$ of these cubes obey $(*)$. Similarly, $\mathbf{labels}_2(u)$ make up at least one third of the total number of labels, thus $(*)$ holds for $\approx 1-\frac{3}{2}\alpha$ of them (and in particular for at least one cube in $\mathbf{labels}_2(u)$, which is all we'll need).

Recall from the construction of the offspring-labels that the cubes $\mathcal{C}_v \in \mathbf{labels}_2(u)$ have $\mathbf{dst}(v) = E_{b_{i+1}}(\mathbf{dst}(u))$. By Definition 9.4, these points are mapped to the exact same variables: $\mathbf{var}_u(\mathbf{dst}(u)) = \mathbf{var}_v(\mathbf{dst}(v))$. As long as there exists one good cube $\mathcal{C}_v \in \mathbf{labels}_2(v)$, we'll have for every $x \in \mathbf{dst}(u)$, $y = E_{b_{i+1}}(x) \in \mathbf{dst}(v)$, and by (2), $\pi_x(\tilde{\mathcal{G}}_v) = m(x)$. We have shown,

$$\forall x \in \mathbf{dst}(u) \quad \pi_x(\mathcal{G}_u) = m(x)\,.$$

We have left to show that

$$\Pr_{x \in \mathbf{manf}(u)} (\pi_x(\mathcal{G}_u) = m(x)) \geq 1 - \alpha$$

Consider the second half of u's offspring, labelled by a label from

$$\mathbf{labels}_1(u) = \bigcup_{\mathtt{x} \in \mathbf{var}_u(\mathbf{dom}_u) \setminus V_u} \mathbf{labels}_{\mathtt{x}}(u)$$

These offspring are actually divided into $|\mathbf{manf}(u)|$ parts, one per each point $x \in \mathbf{manf}(u)$ (with the correspondence $\mathbf{var}_u(x) = \mathtt{x}$). By definition, the offspring v in x's sub-part have $\mathbf{dst}(v) \overset{def}{=} \left\{E_{b_{i+1}}(x)\right\}$. We have shown (recall (2) from before) that $\pi_x(\tilde{\mathcal{G}}_v) = m(x)$ holds for any x such that $y = E_{b_{i+1}}(x) \in \mathbf{dst}(v)$. Hence for every $x \in \mathbf{manf}(u)$, if there is a good label $\mathcal{C}_v \in \mathbf{labels}_{\mathbf{var}_u(x)}(u)$, then $\pi_x(\tilde{\mathcal{G}}_v) = m(x)$.

As shown above, $1-\alpha$ of the cubes in $\mathbf{labels}_1(u)$ are both good, and obey $\pi_{\mathcal{C}_v}(\mathcal{G}_u) = \tilde{\mathcal{G}}_v$. Hence for $1-\alpha$ of the points $x \in \mathbf{manf}(u)$ there must be a good cube $\mathcal{C}_v \in \mathbf{labels}_{\mathbf{var}_u(x)}(u)$ for which $\pi_{\mathcal{C}_v}(\mathcal{G}_u) = \tilde{\mathcal{G}}_v$. In this case,

$$\pi_x(\mathcal{G}_u) = \pi_x(\pi_{\mathcal{C}_v}(\mathcal{G}_u)) = \pi_x(\tilde{\mathcal{G}}_v) = m(x)$$

This establishes,

$$\Pr_{x \in \mathbf{manf}(u)} (\pi_x(\mathcal{G}_u) = m(x)) \geq 1 - \alpha$$

■

10.3 The Consistency Lemma

In this section we prove the consistency lemma, that allows us to deduce one global super-LDF for any good node, assuming "consistent" LDFs on its good offspring.

Lemma 10.5 (Consistency Lemma) *Let $u \in \mathbf{nodes}_i^*$ for some $0 \le i < L$. Define $\mathcal{S}^*$ to be the multi-set of cubes that label good offspring of u, $\mathcal{S}^* \overset{def}{=} \left\{\mathcal{C}_v \in \mathbf{labels}(u) \mid v \in \mathbf{nodes}_{i+1}^*\right\}$. If for every good offspring v of u there is a super-LDF $\tilde{\mathcal{G}}_v$ over $\mathcal{C}_v$, of total degree $\le r = \tilde{r}_i/(D+2)$ and norm $\|\tilde{\mathcal{G}}_v\| \le 2^{L-i-1} \cdot \mathrm{Avg}(v)$, such that*

$$\Pr_{x \in \mathcal{C}_v}\left(\pi_x(\tilde{\mathcal{G}}_v) = m(x)\right) \ge 1 - \alpha$$

then there is a super-LDF $\mathcal{G}_u$ over $\mathbf{dom}_u$ of total degree $\tilde{r}_i = r(D+2)$ and norm $\|\mathcal{G}_u\| \le 2^{L-i} \cdot \mathrm{Avg}(u)$ that obeys

$$\Pr_{\mathcal{C}_v \in \mathcal{S}^*}\left(\pi_{\mathcal{C}_v}(\mathcal{G}_u) = \tilde{\mathcal{G}}_v\right) \ge 1 - \alpha/6$$

Proof: Throughout the following proof, we make no effort to minimize the constants, but rather to shorten the mathematical expressions in which they appear.

Unless otherwise mentioned, $\mathcal{C}_v$ will denote the cube labeling the node v.

For simplicity, assume $\mathbf{dom}_u = \mathcal{F}^d$ ($\mathcal{F}^{d_0}$ in case $i = 0$).

An $[r,d]$-LDF $P : \mathcal{F}^d \to \mathcal{F}$ is called *permissible with coefficient c_P* if $c_P \ne 0$ and for at least $2/3$ of the cubes $\mathcal{C}_v \in \mathcal{S}^*$, $\tilde{\mathcal{G}}_v[P|_{\mathcal{C}_v}] = c_P$. We define the global super-LDF $\mathcal{G}_u$ by

$$\forall P \in LDF_{r,d}: \quad \mathcal{G}_u[P] \overset{def}{=} \begin{cases} c_P & P \text{ is permissible with } c_P \\ 0 & P \text{ isn't permissible} \end{cases}$$

We claim that $\mathcal{G}_u$ is the desired global super-LDF. We first claim that $\|\mathcal{G}_u\| \le 2{\cdot}2^{L-i-1}\mathrm{Avg}(u)$.

Proposition 10.6 *The norm of $\mathcal{G}_u$ is bounded by $2 \cdot 2^{L-i-1}\mathrm{Avg}(u) = 2^{L-i}\mathrm{Avg}(u)$.*

Proof: Denote by $P_1, .., P_a$ the permissible LDFs (if $a = 0$ we're done), and denote $c_i \overset{def}{=} \mathcal{G}_u[P_i]$. Let us consider the average A of the norms $\|\tilde{\mathcal{G}}_v\|$,

$$\begin{aligned} \mathrm{A} &\overset{def}{=} \frac{1}{|\mathcal{S}^*|}\sum_{\mathcal{C}_v \in \mathcal{S}^*} \|\tilde{\mathcal{G}}_v\| \le \frac{1}{|\mathcal{S}^*|}\sum_{\mathcal{C}_v \in \mathcal{S}^*} 2^{L-i-1}\mathrm{Avg}(v) \le 2^{L-i-1}\mathrm{Avg}(u) \\ &\le 2^{L-i-1}g_i = 2^{L-i-1} \cdot g \cdot C^{i+1} < |\mathcal{F}|^{c_g} \cdot C^L \ll \sqrt{|\mathcal{F}|} \end{aligned}$$

where the second inequality in the first line is true since averaging the norm over all of the offspring is at least as large as the average of the good offspring.

We will lower bound A as follows. P_1 appears with c_1 in $\ge 2/3$ of the good cubes, which means $\mathrm{A} \ge 2/3 \cdot |c_1|$. P_2 appears with c_2 in $\ge 2/3$ of the good cubes, however some of its appearances can coincide with those of P_1. Denote by γ the maximal fraction of cubes on which possibly $P_1|_{\mathcal{C}} = P_2|_{\mathcal{C}}$ (by Proposition 8.2, $\gamma \leq \frac{rd}{|\mathcal{F}|}$). Hence A $\geq 2/3 \cdot |c_1| + (2/3 - \gamma) \cdot |c_2|$. Continuing in this manner, P_3 adds at least $(2/3 - 2\gamma) \cdot |c_3|$ and we obtain

$$\forall 1 \leq j \leq a \qquad \mathrm{A} \geq \sum_{i=1}^{j} \left(\frac{2}{3} - (i-1)\gamma \right) \cdot |c_i|$$

If $a \geq \frac{1}{6\gamma} \geq \frac{|\mathcal{F}|}{6rd}$, we get A $\geq \sum_{i=1}^{1/6\gamma} \left(\frac{2}{3} - \frac{1}{6\gamma}\gamma \right) \cdot |c_i| \geq \frac{1}{6\gamma} \cdot \frac{1}{2} \cdot 1 \gg \sqrt{|\mathcal{F}|}$, a contradiction. Thus $a < 1/6\gamma$, and

$$\mathrm{A} \geq \sum_{i=1}^{a} \left(\frac{2}{3} - (i-1)\gamma \right) \cdot |c_i| \geq \frac{1}{2} \sum_{i=1}^{a} |c_i| = \frac{1}{2} \|\mathcal{G}_u\|$$

and indeed $\|\mathcal{G}_u\| \leq 2\mathrm{A} \leq 2 \cdot 2^{L-i-1} \mathrm{Avg}(u) = 2^{L-i} \mathrm{Avg}(u)$. ■

We have left to show that for almost all of the cubes $\mathcal{C}_v \in \mathcal{S}^*$, $\pi_{\mathcal{C}_v}(\mathcal{G}_u) = \tilde{\mathcal{G}}_v$.

Let us define, for every good node v, the remainder super-LDF: $\mathcal{R}_v \stackrel{def}{=} \tilde{\mathcal{G}}_v - \pi_{\mathcal{C}_v}(\mathcal{G}_u)$ (the definition of $\mathcal{G}_u$ implies that every LDF P appearing in it has degree $\leq r$; subtraction is defined as usual subtraction of two vectors in $\mathbb{Z}^{LDF_{r,D+2}}$). Assume, for contradiction, that for at least an $\alpha/6$ fraction of the good nodes, $\mathcal{R}_v \neq \vec{0}$. We will derive a contradiction by finding an LDF P that appears with the same coefficient $c_P \neq 0$ in $\mathcal{R}_v$ in at least $2/3 + \gamma\|\mathcal{G}_u\|$ fraction of the good nodes v. This LDF P can agree with another LDF in $\mathcal{G}_u$ on at most $\gamma\|\mathcal{G}_u\|$ fraction of the good cubes. Hence on at least 2/3 of the good cubes, $c'_P \stackrel{def}{=} \tilde{\mathcal{G}}_v[P|_{\mathcal{C}_v}] = c_P + \mathcal{G}_u[P]$, which implies that P is permissible with c'_P, so by definition $\mathcal{G}_u[P] = c'_P$, hence $c_P = 0$, a contradiction.

For every $x \in \mathcal{F}^d$, define $m_R(x) \stackrel{def}{=} m(x) - \pi_x(\mathcal{G}_u)$. Obviously $m_R(x) = \pi_x(\mathcal{R}_v)$ if and only if $m(x) = \pi_x(\tilde{\mathcal{G}}_v)$. (This happens for at least $1 - \alpha$ of the points $x \in \mathcal{C}_v$ for every $\mathcal{C}_v \in \mathcal{S}^*$, by the conditions of the lemma).

Proposition 10.7 *Let $\mathcal{R}_v \stackrel{def}{=} \tilde{\mathcal{G}}_v - \pi_{\mathcal{C}_v}(\mathcal{G}_u)$ be as before. There exists an $[r, D+2]$-LDF P and a coefficient $c_P \neq 0$ such that*

$$\Pr_{\mathcal{C}_v \in \mathcal{S}^*} \left(\mathcal{R}_v[P|_{\mathcal{C}_v}] = c_P \right) > \delta$$

where $\delta = \Omega\left((\frac{\alpha}{s})^9 \right)$ and $s \stackrel{def}{=} 2^{L-i}\mathrm{Avg}(u)$.

Proof: Consider the following random procedure:

1. For every cube $\mathcal{C}_v \in \mathbf{labels}(u)$ choose a random LDF from the set $\{ Q \in LDF_{r,D+2} \mid \mathcal{R}_v[Q] \neq 0 \}$. If $\mathcal{C}_v \notin \mathcal{S}^*$ or this set is empty, choose nothing.
2. For every point $x \in \mathbf{dom}_u$ choose a random value from the set $\{ a \in \mathcal{F} \mid m_R(x)[a] \neq 0 \}$. If this set is empty, choose nothing.

3. Choose a random cube $\mathcal{C}_v \in \mathbf{labels}(u)$ and a random point $x \in \mathcal{C}_v$. If no value is chosen for either the point or the cube, the procedure fails.

We are interested in pairs of good cube and point on it, on which the procedure doesn't fail, and that have relatively few possible values to choose from, and that are consistent. We eliminate 'bad' pairs as follows.

For a cube $\mathcal{C}_v \in \mathbf{labels}(u)$ define $E_1(\mathcal{C}_v)$ to be the predicate that evaluates to `true` iff $\mathcal{C}_v \in \mathcal{S}^*$ and the set $\{Q \in LDF_{r,D+2} \mid \mathcal{R}_v[Q] \neq 0\}$ is non-empty. $\Pr_{\mathcal{C}_v \in \mathbf{labels}(u)}(E_1(\mathcal{C}_v)) \geq (1-1/C) \cdot \alpha/6$ because $1-1/C$ of the cubes are in $\mathcal{S}^*$ (since u is good), and we assumed for contradiction that for $\alpha/6$ of these cubes $\mathcal{R}_v$ is non-trivial.

For a cube $\mathcal{C}_v \in \mathbf{labels}(u)$ define $E_2(\mathcal{C}_v)$ to be the predicate that evaluates to `true` iff $E_1(\mathcal{C}_v)$ is true and also $\|\tilde{\mathcal{G}}_v\| \leq 2 \cdot \frac{6s}{\alpha(1-1/C)}$ where $s = 2^{L-i}\text{Avg}(u)$ bounds the average of $\|\tilde{\mathcal{G}}_v\|$ taken over nodes $v \in \mathcal{S}^*$. We note that the average norm $\|\tilde{\mathcal{G}}_v\|$ taken over cubes for which E_1 is true does not exceed $\frac{6s}{\alpha(1-1/C)}$. The standard Markov argument shows

$$\Pr_{\mathcal{C}_v \in \mathbf{labels}(u)}(E_2(\mathcal{C}_v)) \geq \frac{1}{2} \cdot \Pr_{\mathcal{C}_v \in \mathbf{labels}(u)}(E_1(\mathcal{C}_v)) \geq (1-1/C) \cdot \alpha/12$$

By the triangle inequality, $\|\mathcal{R}_v\| = \|\tilde{\mathcal{G}}_v - \pi_{\mathcal{C}_v}(\mathcal{G}_u)\| \leq \|\tilde{\mathcal{G}}_v\| + \|\mathcal{G}_u\| \leq \|\tilde{\mathcal{G}}_v\| + s$ hence the cubes $\mathcal{C}_v$ for which $E_2(\mathcal{C}_v) = \texttt{true}$ have $\|\mathcal{R}_v\| \leq 12s/\alpha(1-1/C) + s < 13s/\alpha$.

For a cube $\mathcal{C}_v \in \mathbf{labels}(u)$, and a point $x \in \mathcal{C}_v$, define $E_3(\mathcal{C}_v, x)$ to be the predicate that evaluates to `true` iff $E_2(\mathcal{C}_v) = \texttt{true}$ and also $x \in \mathcal{C}_v$ and x obeys $m_R(x) = \pi_x(\mathcal{R}_v)$ *and* $\mathcal{R}_v$ is not ambiguous on x. We say, in this case, that the point x and the cube $\mathcal{C}_v$ *agree non-ambiguously.* Since for every good cube $\mathcal{C}_v$ no more than α fraction of its points have $m_R(x) \neq \pi_x(\mathcal{R}_v)$, and no more than $\text{amb}(r, D+2, \|\mathcal{R}_v\|)$ are ambiguous, it follows that

$$\begin{aligned}
\Pr_{\substack{\mathcal{C}_v \in \mathbf{labels}(u) \\ x \in \mathcal{C}_v}}(E_3(\mathcal{C}_v, x)) &\geq \Pr_{\mathcal{C}_v}(E_2(\mathcal{C}_v)) \cdot (1 - \alpha - \text{amb}(r, D+2, \|\mathcal{R}_v\|)) \\
&\geq (1-1/C) \cdot \alpha/12 \cdot (1 - \alpha - |\mathcal{F}|^{-\frac{1}{2}}) \\
&> \alpha/100
\end{aligned}$$

(the second inequality follows from $\text{amb}(r, D+2, \|\mathcal{R}_v\|) \leq \frac{\tilde{r}_i/(D+2)\cdot(D+2)}{|\mathcal{F}|} \cdot \|\mathcal{R}_v\|^2 \ll |\mathcal{F}|^{-\frac{1}{2}}$). The pairs of point x and cube $\mathcal{C}_v$ for which $E_3(\mathcal{C}_v, x) = \texttt{true}$ are pairs that agree non-ambiguously, and for which $\|\mathcal{R}_v\| \leq 13s/\alpha$.

For a cube $\mathcal{C}_v \in \mathbf{labels}(u)$, a point $x \in \mathcal{C}_v$, an $[r, D+2]$-LDF Q (viewed as an LDF over $\mathcal{C}_v$) and a value $a \in \mathcal{F}$, define $E_4(\mathcal{C}_v, x, Q, a)$ to be the predicate that evaluates to `true` iff $E_3(\mathcal{C}_v, x) = \texttt{true}$ and also $Q(x) = a$. We will lower bound the probability $\Pr_{\mathcal{C}_v, x, Q, a}(E_4(\mathcal{C}_v, x, Q, a))$ where $\mathcal{C}_v \in \mathbf{labels}(u)$, $x \in \mathcal{C}_v$, and Q and a are chosen according to the random procedure described in the beginning of the proof (i.e. Q is chosen uniformly from the set $\{Q \in LDF_{r,D+2} \mid \mathcal{R}_v[Q] \neq 0\}$, and a uniformly from the set $\{a \in \mathcal{F} \mid m_R(x)[a] \neq 0\}$. Note that when E_4 is true, there are no more than $\frac{13s}{\alpha}$ LDFs that appear in $\mathcal{R}_v$. Since $E_4(\mathcal{C}_v, x, Q, a) = \texttt{true}$ implies by definition $E_3(\mathcal{C}_v, x) = \texttt{true}$ we know that $m_R(x) = \pi_x(\mathcal{R}_v)$, hence any value a randomly chosen for x has a "matching value" in the set $\{Q \in LDF_{r,D+2} \,|\, \mathcal{R}_v[Q] \neq 0\}$. This value is the chosen one with probability at least $\frac{\alpha}{13s}$. Thus,

$$\Pr(E_4(\mathcal{C}_v, x, Q, a)) \geq \Pr(E_3(\mathcal{C}_v, x)) \cdot \frac{\alpha}{13s}$$

Finally, note that if $E_4(\mathcal{C}_v, x, Q, a) = \texttt{true}$, then $\mathcal{R}_v[Q] = m_R(x)[a]$ because the cube and point agree non-ambiguously. Also, since in this case $\|\mathcal{R}_v\| \leq 13s/\alpha$, the coefficient $\mathcal{R}_v[Q]$ can be any value from the set $B \overset{def}{=} \{\pm 1, ..., \pm 13s/\alpha\}$, $26s/\alpha$ values in all. Denote by $E_c(\mathcal{C}_v, x, Q, a)$ the predicate that is the same as E_4 except that it evaluates to true only if in addition, $\mathcal{R}_v[Q] = m_R(x)[a] = c$. There must be at least one value $c_0 \in B$ for which

$$\Pr(E_{c_0}(\mathcal{C}_v, x, Q, a)) \geq \frac{\alpha}{26s} \cdot \Pr(E_4(\mathcal{C}_v, x, Q, a)) \geq \frac{\alpha}{26s} \cdot \frac{\alpha}{13s} \cdot \Pr(E_3(\mathcal{C}_v, x)) \geq \frac{\alpha^2}{338s^2} \cdot \frac{\alpha}{100} = \Omega(\frac{\alpha}{s})^3$$

We now apply the following corollary of [RS97],

Lemma 10.8 *Let $\rho = (\frac{rd}{\mathcal{F}})^c$ for some constant $c > 0$, and let $\mathcal{S} = \mathbf{labels}(u)$ for labels(u) as above. Let $\mathcal{A} : \mathcal{S} \to \mathrm{LDF}_{r,D+2}$ be an assignment of $[r, D+2]$-LDF per cube, and let $\mathcal{A}_0 : \mathcal{F}^d \to \mathcal{F}$ be an assignment of value per point. If*

$$\Pr_{\mathcal{C} \in_R \mathcal{S}, x \in_R \mathcal{C}} (\mathcal{A}[\mathcal{C}](x) = \mathcal{A}_0[x]) \geq \rho$$

then there is an $[r, d]$-LDF P for which $\Pr_{\mathcal{C} \in \mathcal{S}}(P|_{\mathcal{C}} = \mathcal{A}[\mathcal{C}]) \geq \rho^3$.

We omit the proof of this lemma, and note that a very similar cube vs. point version appears in [DFK+99]. We apply this lemma as follows. We take $\mathcal{S} = \mathbf{labels}(u)$. For every cube $\mathcal{C}_v$ whose selected value Q has $\mathcal{R}_v[Q] = c_0$, assign $\mathcal{A}[\mathcal{C}_v] = Q$, otherwise let $\mathcal{A}[\mathcal{C}_v]$ be a totally random value. For each point $x \in \mathcal{F}^d$, we define $\mathcal{A}_0[x]$ to be the value selected for it in the random procedure. Again, if no value was selected, we assign a totally random value. We have

$$\Pr_{\mathcal{C} \in \mathcal{S}, x \in \mathcal{C}} (\mathcal{A}[\mathcal{C}](x) = \mathcal{A}_0[x]) \geq \Pr(E_{c_0}(\mathcal{C}, x, \mathcal{A}[\mathcal{C}], \mathcal{A}_0[x]))$$

The probability on the right hand side is taken over a random choice of cube $\mathcal{C} \in \mathcal{S}$ and point $x \in \mathcal{C}$, and over the random choices made when defining $\mathcal{A}[\mathcal{C}]$ and $\mathcal{A}_0[x]$. It follows easily that this probability is at least $\geq \Pr(E_{c_0})$. Thus we obtain an *LDF* P that agrees with $\geq (\Pr(E_{c_0}))^3 \geq \Omega\left((\frac{\alpha}{s})^9\right)$ fraction of the cubes and their chosen values. P expectedly appears in less than $1/|\mathcal{F}|$ of the cubes that were randomly assigned. Thus at least half of the cubes in which P appears also obey $\mathcal{R}_v[P|_{\mathcal{C}_v}] = c_0$. These cubes make up at least $\delta = \frac{1}{2} \cdot (\Pr(E_{c_0}))^3 = \Omega\left((\frac{\alpha}{s})^9\right)$ of the good cubes. ■

We have found a polynomial P that appears (with the same coefficient $c_0 \neq 0$) in a non-negligible fraction of the cubes (i.e. $\mathcal{R}_v[P|_{\mathcal{C}_v}] = c_0$ for a non-negligible $\geq \delta$ fraction of the good offspring v of u). We now show that P, in fact, appears in most of the points with coefficient c_0,

Proposition 10.9 *For most points $x \in \mathbf{dom}_u$, $m_R(x)[P(x)] = c_0$.*

Proof: Let $N = \{x \in \mathbf{dom}_u \mid m_R(x)[P(x)] \neq c_0\}$ be the set of points where P does not appear with coefficient c_0. We shall prove that $\mu \overset{def}{=} \frac{|N|}{|\mathbf{dom}_u|} < \frac{1}{2}$. We now state a hitting lemma that shows that if N is not too small, then almost all of the cubes must hit a non-negligible fraction of the points in N.

Lemma 10.10 (Hitting Lemma) *Let $0 < \beta < 1$ and let $\mathcal{D} = \mathcal{F}^d$. Let $N \subset \mathcal{D}$ be a set of points, $|N| \geq \beta\,|\mathcal{D}|$. Most $(1 - \frac{8}{\beta|\mathcal{F}|})$ cubes in $\mathbf{labels}(u)$ (for u as above) have at least $\frac{\beta}{2}$ of their points in N.*

The proof of this lemma is easily obtained using the pairwise independence of points in a random cube, and is omitted (special care should be given to the fact that the points in these cubes are distributed only *almost* uniformly: certain points – e.g. span($\mathbf{dst}(u)$) – appear more often than others).

We now know that $1 - \frac{8}{\mu|\mathcal{F}|}$ of the cubes in $\mathbf{labels}(u)$ $(1 - \frac{16}{\mu|\mathcal{F}|}$ fraction of $\mathcal{S}^*$, since $|\mathcal{S}^*| > \frac{1}{2}\,|\mathbf{labels}(u)|)$ have $\frac{\mu}{2}$ of their points from N. Consider only cubes $\mathcal{C}_v$ whose norm isn't too large – $\|\mathcal{R}_v\| \leq 2s/(\delta/2)$ (the average of $\|\mathcal{R}_v\|$ over all nodes $v \in \mathcal{S}^*$ is $\leq s{+}s = 2s$, hence we are ignoring a $\delta/2$ fraction). If $\frac{\mu}{2} > \alpha + \mathrm{amb}(r, D+2, 2s/(\delta/2))$ then every such cube must agree non-ambiguously with at least one point from N. This implies that P does not appear in these cubes (that constitute at least $1 - \frac{16}{\mu|\mathcal{F}|} - \delta/2$ fraction of $\mathcal{S}^*$) with coefficient c_0, and hence, $\delta \leq \frac{16}{\mu|\mathcal{F}|} + \delta/2$. Altogether we have that

$$\mu \leq \max\left(2(\mathrm{amb}(r, D+2, 2s/(\delta/2)) + \alpha), \frac{32}{\delta\,|\mathcal{F}|}\right) < \frac{1}{2}$$

■

Having P appearing in most points, we now show that P appears in most cubes with coefficient c_0.

Proposition 10.11 *For at least $3/4$ of the cubes $\mathcal{C}_v \in \mathcal{S}^*$, $\mathcal{R}_v[P|_{\mathcal{C}_v}] = c_0$.*

Proof: Let $N = \{x \in \mathbf{dom}_u \mid m_R(x)[P(x)] = c_0\}$. N has, by Proposition 10.9, most of the points in $\mathcal{F}^d$. According to the hitting lemma, all except $\frac{16}{|\mathcal{F}|}$ of the cubes in $\mathbf{labels}(u)$ $(\frac{32}{|\mathcal{F}|}$ of $\mathcal{S}^*)$, have $\frac{1}{4}$ of their points from N.

By the Markov inequality, at most $1/10$ of the cubes in $\mathcal{S}^*$ have norm $\|\mathcal{R}_v\| \leq 10 \cdot 2s = 20s$, and thus no more than $20s$ LDFs appearing in them. Therefore $1 - 1/10 - \frac{32}{|\mathcal{F}|} > 3/4$

of the cubes in $\mathcal{S}^*$ have $\frac{1}{4}$ of their points from N, and are assigned no more than $20s$ LDFs. Denote these cubes $\mathcal{S}^*(P)$. We will show that for every cube $\mathcal{C}_v \in \mathcal{S}^*(P)$, $\mathcal{R}_v[P|_{\mathcal{C}_v}] = c_0$.

Let $\mathcal{C}_v$ be a cube in $\mathcal{S}^*(P)$. The fraction of points of $\mathcal{C}_v$ on which m_R agrees with $\mathcal{R}_v$ non-ambiguously and the point belongs to N is at least $\frac{1}{4} - \alpha - \mathrm{amb}(r, D+2, 20s) > \frac{1}{5}$ (recall $\alpha \leq 1/100$). For each such point $x \in \mathcal{C}_v$, there is an LDF Q, $Q(x) = P(x)$ with $\mathcal{R}_v[Q] = c_0$. For every such point there are no more than $20s$ candidates, hence there is at least one LDF Q with $\mathcal{R}_v[Q] = c_0$ that is equal to P on at least

$$\frac{1}{5} \cdot \frac{1}{20s} > \frac{r(D+2)}{|\mathcal{F}|}$$

of $\mathcal{C}_v$'s points. This LDF is therefore equal to $P|_{\mathcal{C}_v}$ (two distinct $[r, D+2]$-LDFs can agree on at most $\frac{r(D+2)}{|\mathcal{F}|}$ fraction of their domain).

We have shown that $\mathcal{R}_v[P|_{\mathcal{C}_v}] = c_0$ for all cubes $\mathcal{C}_v \in \mathcal{S}^*(P)$, which make up at least $3/4$ of the cubes in $\mathcal{S}^*$. ■

We unveiled an LDF P that appears with the same coefficient $c_0 \neq 0$ in $\mathcal{R}_v$ for at least $3/4 > 2/3 + \gamma\|\mathcal{G}_u\|$ of the good nodes v. Hence P appears with the same ($c' = c_0 + \mathcal{G}_u[P]$) coefficient in $\tilde{\mathcal{G}}_v$ for at least $2/3$ of the good nodes v. Thus, P is permissible with coefficient c', and by our definition of $\mathcal{G}_u$, $\mathcal{G}_u[P] = c'$. Thus $c_0 = 0$, a contradiction. ■

Chapter 11

g-CVP is NP-hard

We begin by defining the Closest Vector Problem (CVP), and its gap version g-CVP. We then define an intermediate problem called Shortest Integer Solution (SIS), and show a reduction from g-SIS to g-CVP. We then show the simple reduction from g-S-SAT to g-SIS and therefore to g-CVP. We restrict ourselves to l_1 norm, although the results can be easily translated to any l_p norm, $1 \leq p < \infty$.

A lattice $L = L(v_1, .., v_n)$, for linearly independent vectors $v_1, .., v_n \in R^k$ is the set of all integral linear combinations of $v_1, .., v_n$, $L = \{\sum a_i v_i \mid a_i \in \mathbb{Z}\}$.
The closest-vector problem is defined as follows:

CVP. Given (L, y) where $L = L(v_1, .., v_n)$ is a lattice and $y \in R^k$, find a lattice vector closest to y (i.e. a lattice vector $v \in L$ that minimizes $\|v - y\|$.

Approximating CVP to within factor $g = g(n)$ means finding a lattice vector v whose distance from y, $\|v - y\|$, is no more than g times the minimal distance. The gap version of CVP is a decision problem as follows,

g-CVP. Given (L, y, d) for a lattice L, a vector $y \in R^k$, and a number d, distinguish between the following two cases:

Yes: There exists a lattice vector $v \in L$ for which $\|v - y\| \leq d$.

No: For every lattice vector $v \in L$, $\|v - y\| > g \cdot d$.

Proving that g-CVP is NP-hard means that having an approximation algorithm to within factor g were to imply $P = NP$.

11.1 Shortest Integer Solution - SIS

Definition of SIS and g-SIS

We define a variant of CVP named Shortest Integer Solution (SIS) and its gap version, g-SIS. We then show a simple reduction from g-SIS to g-CVP.

SIS: Given (B, t) for an integer matrix B with columns $b_1, .., b_n$ and a target vector $t \in L(b_1, ..., b_n)$, find integer coefficients a_i such that $\sum a_i b_i = t$ (we assume such a_i exist), and such that the *length* $\sum |a_i|$ of the solution is minimal. In other words, find the shortest integer solution for the linear system $B \cdot x = t$.

The gap version of SIS is as follows,

g-**SIS**: Given (B, t, d) with B and t as before, and a number d, distinguish between the following two cases:

Yes: The shortest integer solution is of length d or less.

No: The shortest integer solution is of length $> g \cdot d$.

Reducing g-SIS to g-CVP

Given an instance of g-SIS, (B, t, d), we efficiently construct a lattice L and a target vector y such that 'yes' instances of g-SIS are translated into 'yes' instances of g-CVP and 'no' instances are translated into 'no' instances. The lattice L is constructed by multiplying the matrix B by a very large number w, and adding a distinct 1-coordinate to each column. The vector y (that we are to approximate from within the lattice) will be t multiplied by w with zeros in the n additional coordinates:

$$L = \begin{pmatrix} & wB & \\ & & \\ 1 & & 0 \\ & \ddots & \\ 0 & & 1 \end{pmatrix} \qquad y = \begin{pmatrix} \vdots \\ wt \\ \vdots \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

To see that 'yes' instances map into 'yes' instances just note that any solution a, $B \cdot a = t$, gives a lattice vector $L \cdot a$ such that $\|L \cdot a - y\| = \|a\|$. Let w be such that the entries in the upper half of the matrix are all integer multiples of $g \cdot d + 1$. The next lemma will show that 'no' instances of g-SIS (where the shortest solution is of length $> g \cdot d$) map into 'no' instances of g-CVP.

Lemma 11.1 *If there is a lattice vector, $L \cdot a$, such that $r \overset{def}{=} \|L \cdot a - y\| \leq g \cdot d$, then there is an integer solution to (B, t) of length r.*

Proof: $r \leq gd$ means that $L \cdot a = y$ in all but the lower n coordinates, otherwise the distance r would be at least $g \cdot d + 1$. In other words, a is a solution to the g-SIS instance. The lower n coordinates of $L \cdot a$ are exactly equal to a, and therefore $\|a\| = r$. ■

11.2 From S-SAT to g-SIS

We shall prove that g-SIS is NP hard for $g = n^{c/\log\log n}$ (for some constant $c > 0$) by reducing g-S-SAT to it.

We begin with a g-S-SAT test system $I = \langle \Psi = \{\psi_1, .., \psi_n\}, V = \{v_1, .., v_m\}, \{\mathcal{R}_{\psi_1}, .., \mathcal{R}_{\psi_n}\}\rangle$ where Ψ is a set of tests over variables V, and for each $\psi \in \Psi$, $\mathcal{R}_\psi$ is the set of satisfying assignments for ψ. We (efficiently) construct from it an instance of g-SIS, (B, t, d). We then show that the 'yes' instances of g-S-SAT are mapped to 'yes' instanceof g-SIS and 'no' instances to 'no' instances.

We show that a consistent natural super-assignment to Ψ translates to a short (i.e. of l_1 norm $|\Psi|$) solution for (B, t). On the other hand we show that any solution that is shorter than $g \cdot |\Psi|$, translates to a consistent super-assignment of norm $< g$ for Ψ.

The General Construction. The matrix B will have a column for every pair of test $\psi \in \Psi$ and a satisfying assignment $r \in \mathcal{R}_\psi$ for it. The upper rows of B will take care of consistency, and the lower rows will take care of non-triviality.

Non-Triviality Rows. There will be a row designated to each test. In the row of ψ all of the columns associated with ψ will have a 1, and all other columns will have zero.

Consistency Rows. We shall have $|\mathcal{F}|$ rows for each pair of tests ψ_i and ψ_j and common variable x (there will be $a \cdot |\mathcal{F}|$ rows if ψ_i and ψ_j share a variables). These rows serve as a consistency-ensuring gadget and only the columns associated with ψ_i and ψ_j will have non-zero values in these rows. The gadget will ensure that the super-assignments to ψ_i and ψ_j *are consistent* on their common variable x.

The **target vector** t will be an all-1 vector. We set $d \overset{def}{=} |\Psi|$.
We now turn to describe the structure of the gadget itself. This will complete the description of the g-SIS instance.

The Gadget. Let's concentrate on the gadget for the pair of tests ψ_i and ψ_j with common variable x. This is a pair of matrices G_1 of dimension $(|\mathcal{F}| \times |\mathcal{R}_{\psi_i}|)$ and G_2 of dimension $(|\mathcal{F}| \times |\mathcal{R}_{\psi_j}|)$. The matrices G_1 and G_2 have $|\mathcal{F}|$ rows, each corresponding to a possible assignment for the variable x. The r-th column in G_1 is the 'characteristic function' of $r|_x$, i.e. zeros everywhere except for a 1 in the $r|_x$-th coordinate. Similarly, the column in G_2

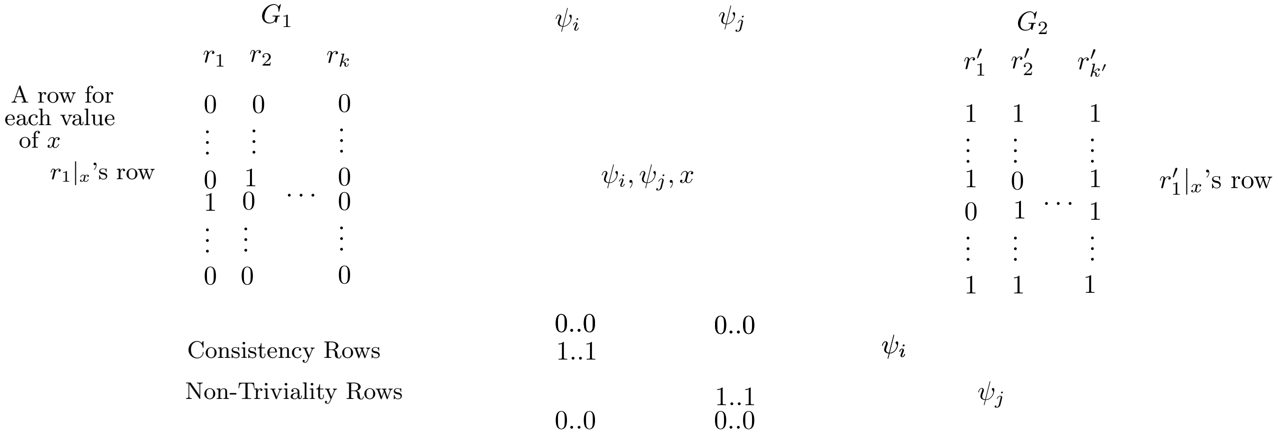

Figure 11.1: The SIS matrix B

corresponding to r' is the negation of the characteristic function of $r'|_x$, i.e. 1 everywhere except for one 0 in the $r'|_x$-th coordinate (see Figure 11.1).

Proving Correctness. Let us now show that 'yes' instances of the S-SAT map to 'yes' instances of the g-SIS.

Lemma 11.2 *If there is a consistent natural super-assignment to the* S-SAT *test system* Ψ*, then there is a solution of* l_1 *norm* $|\Psi|$ *to the above g-SIS instance.*

Proof: We take the consistent natural super-assignment S and construct from it a solution to the g-SIS. We will concatenate the vectors $S(\psi_1)S(\psi_2)$... (turning n $|\mathcal{R}_{\psi_i}|$-coordinate vectors into one long vector with $\sum_i |\mathcal{R}_{\psi_i}|$ coordinates) to obtain our alleged solution to g-SIS. The target vector is reached in the non-triviality rows because S is natural i.e. it assigns a $+1$ coefficient to exactly one column of every test.

To show that the target vector is reached in the consistency rows, consider the set of $|\mathcal{F}|$ rows belonging to an arbitrary pair of tests ψ_i and ψ_j with common variable x. Suppose $S(\psi_i)[r_1], S(\psi_j)[r_2]$ are the single 1's in $S(\psi_i), S(\psi_j)$ respectively (S is natural). S is consistent so $r_1|_x = r_2|_x$. By the construction of B we see that

$$\begin{array}{r} \\ 1 \\ \vdots \\ r_1|_x \\ \vdots \\ |\mathcal{F}| \end{array}\begin{array}{c} r_1 \\ \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \end{array} + \begin{array}{r} \\ 1 \\ \vdots \\ r_2|_x \\ \vdots \\ |\mathcal{F}| \end{array}\begin{array}{c} r_2 \\ \begin{pmatrix} 1 \\ 1 \\ 0 \\ 1 \\ 1 \end{pmatrix} \end{array} = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}$$

and the target vector is reached in these rows.

The length of the solution is the sum of the lengths of the $S(x)$'s, and since $\|S\| = 1$, it is exactly $|\Psi|$. ■

We will now show that 'no' instances of the S-SAT map to 'no' instances of the g-SIS by showing that if we ended up with an instance that isn't a 'no' instance, then we must have started with a non-'no' instance.

Lemma 11.3 *Let s be a solution to the above g-SIS instance, $\|s\| \leq g\,|\Psi|$. There exists a non-trivial consistent super-assignment S of norm $\leq g$ for the* S-SAT *instance.*

Proof: We show how to construct S from s: we 'break' s into $|\Psi|$ pieces of length $|\mathcal{R}_\psi|$, one for each test $\psi \in \Psi$. We obtain a super-assignment S whose norm is $\frac{1}{|\Psi|}\|s\|$.

For any arbitrary $\psi \in \Psi$, the target vector is reached in the ψ-th row of the non-triviality rows. This implies that

$$\sum_{r \in \mathcal{R}_\psi} S(\psi)[r] = 1 \tag{11.1}$$

and in particular S is non-trivial (the sum of the coordinates in $S(\psi)$ remains the same under projection to any single variable).

Let $\psi_i, \psi_j \in \Psi$ be arbitrary tests with a common variable x. We shall show that $\pi_x(S(\psi_i)) = \pi_x(S(\psi_j))$. Consider the $|\mathcal{F}|$ rows that correspond to ψ_i, ψ_j, x. In each of these rows the sum of the vectors is 1, in other words, for any $f \in \mathcal{F}$,

$$\sum_{r\,:\,r|_x = f} S(\psi_i)[r] + \sum_{r\,:\,r|_x \neq f} S(\psi_j)[r] = 1 \tag{11.2}$$

Subtracting (11.1) for ψ_j from (11.2) gives,

$$\sum_{r\,:\,r|_x = f} S(\psi_i)[r] = \sum_{r\,:\,r|_x = f} S(\psi_j)[r]$$

which, by definition of the projection means $\pi_x(S(\psi_i)) = \pi_x(S(\psi_j))$. We hence have a consistent super-assignment of norm $\frac{1}{|\Psi|}\|s\| \leq g$.

■

The two above lemmas complete the reduction of S-SAT to g-SIS.

11.3 Other l_p norms

Our result actually holds for CVP with any l_p norm for $1 < p < \infty$, as seen by the following reduction.

Let us begin by observing that in our reduction from S-SAT to CVP via g-SIS, a 'yes' instance (a test-system with a consistent natural super-assignment), was transformed to a g-SIS instance having a solution of length $|\Psi|$, which was transformed to a CVP instance $(L, y, dist = |\Psi|)$ such that there is a lattice vector $v \in L$ with $\|v - y\|_1 = |\Psi|$, and such that the vector $v - y$ is a zero-one vector, thus $\|v - y\|_p = \sqrt[p]{\|v - y\|_1} = \sqrt[p]{|\Psi|}$.

Now take the same lattice L and target vector y as a CVP$_p$ instance with distance parameter $\sqrt[p]{|\Psi|}$, $(L, y, dist_p = \sqrt[p]{|\Psi|})$. The above observation simply says that a 'yes' instance has a solution whose distance is $\sqrt[p]{|\Psi|}$.

On the other hand, if $(L, y, |\Psi|)$ is a 'no' CVP$_1$ instance, then every lattice vector $v \in L$, has $\|v - y\|_1 > g \cdot |\Psi|$. Since $\|x\|_p \geq \sqrt[p]{\|x\|_1}$ for any integer-vector x, we have $\|v - y\|_p > \sqrt[p]{g \cdot |\Psi|} = \sqrt[p]{g} \cdot \sqrt[p]{|\Psi|}$.

This establishes that it is NP-hard to approximate CVP$_p$ to within a factor of $\sqrt[p]{g} = n^{c_p / \log\log n}$ for some constant $c_p > 0$.

Chapter 12

Discussion

Our result for the Closest Vector Problem was obtained via S-SAT using recursive composition, that alternates between two types of algebraic encodings: the embedding extension, and the low-degree extension. This technique was adapted from the proof of a low error-probability PCP characterization of NP [DFK$^+$99], and proved to be useful in this setting as well.

Two interesting open problems remain. The first is the Shortest Vector Problem, the homogeneous counterpart of CVP. This problem is easier to approximate than CVP, as an approximation algorithm for CVP yields an approximation for SVP [GMSS99], yet currently the best approximation algorithms for it give no better factors than those for CVP. However, where hardness results go, the SVP lags behind, with known hardness of approximation for a factor no larger than some constant.

The hardness of approximating SVP is of special interest in cryptography, where the hardness of this problem serves as the basic assumption of a crypto-system of Ajtai and Dwork, see [AD97].

The second open problem is to achieve hardness of approximation factors for CVP that are polynomial in n, say n^ϵ for some $\epsilon > 0$. Our technique seems incapable of doing this, as the recursive structure requires super-constant depth, limiting the blow-up allowable at each level.

Appendix A

Weighted vs Unweighted

We show a simple reduction from Vertex-Cover on weighted graphs to the non-weighted case.

Lemma A.1 *Let $\varrho > 0$ be an arbitrary precision parameter. Given a weighted graph $G = (V, E, \Lambda)$, one can construct, in polynomial time in $|G|, \frac{1}{\varrho}$ an unweighted graph $G_\varrho = (V_\varrho, E_\varrho)$ such that*

$$\left| \frac{\overline{\alpha}(G_\varrho)}{|V_\varrho|} - \overline{\alpha}(G) \right| \leq \varrho$$

Proof: Let $m = |V| \cdot \frac{1}{\varrho}$. We replace each $v \in V$ with $m_v = \lceil m \cdot \Lambda(v) \rfloor$ copies ($\lceil x \rfloor$ denotes the integer nearest x), and set

$$\begin{aligned} V_\varrho &\overset{def}{=} \{\langle v, i\rangle \mid v \in V,\ 1 \leq i \leq m_v\} \\ E_\varrho &\overset{def}{=} \{\{\langle v_1, i_1\rangle, \langle v_2, i_2\rangle\} \mid \{v_1, v_2\} \in E,\ i_1 \in [m_{v_1}],\ i_2 \in [m_{v_2}]\} \end{aligned}$$

If $C \subseteq V$ is a vertex cover for G, then $C_\varrho = \bigcup_{v \in C} \{v\} \times [m_v]$ is a vertex cover for G_ϱ. Moreover, every minimal vertex cover $C_\varrho \subseteq V_\varrho$ is of this form. Thus we show $\left| \frac{\overline{\alpha}(G_\varrho)}{|V_\varrho|} - \overline{\alpha}(G) \right| \leq \varrho$ by the following proposition,

Proposition A.2 *Let $C \subseteq V$, and let $C_\varrho = \bigcup_{v \in C} \{v\} \times [m_v]$. Then $\left| \frac{|C_\varrho|}{|V_\varrho|} - \Lambda(C) \right| \leq \varrho$.*

Proof: For every C, C_ϱ as above, $|C_\varrho| = \sum_{v \in C} m_v = \sum_{v \in C} \lceil m \cdot \Lambda(v) \rfloor = m \cdot \Lambda(C) + \sum_{v \in C} (\lceil m \cdot \Lambda(v) \rfloor - m \cdot \Lambda(v))$. For any z, $|\lceil z \rfloor - z| \leq \frac{1}{2}$, and so

$$\left| \frac{|C_\varrho|}{m} - \Lambda(C) \right| \leq \frac{1}{2} \frac{|C|}{m} \leq \frac{\varrho}{2} \qquad (*)$$

To complete our proof we need to replace $\frac{|C_\varrho|}{m}$ by $\frac{|C_\varrho|}{|V_\varrho|}$ in $(*)$. Indeed, taking $C = V$ in $(*)$, yields $\left| \frac{|V_\varrho|}{m} - 1 \right| \leq \frac{\varrho}{2}$, and multiplying by $\frac{|C_\varrho|}{|V_\varrho|} \leq 1$, we obtain $\left| \frac{|C_\varrho|}{m} - \frac{|C_\varrho|}{|V_\varrho|} \right| \leq \frac{\varrho}{2}$. $\blacksquare$ $\blacksquare$

Appendix B

A Chernoff Bound

In this appendix we make use of the following Chernoff bound [MR97, p. 70],

Theorem B.1 (A Chernoff Bound) *Let $X_1, .., X_l$ be independent Bernoulli variables, s.t. $\forall i$, $\Pr[X_i = 1] = p$. Then for $X = \sum_{i=1}^{l} X_i$, $\mu = E[X] = p \cdot l$ and for any $0 < \delta < 1$,*

$$\Pr[X < (1-\delta)\mu] < exp(-\mu\delta^2/2)$$

This bound directly implies the following proposition,

Proposition B.2 *Let $A : Z \to \{\mathsf{T}, \mathsf{F}\}$ be such that $\Pr_{z \in Z}[A(z) = \mathsf{T}] = \frac{1}{|R_X|}$, then,*

$$\Pr_{B \in \mathcal{Z}}\left[\left|A^{-1}(\mathsf{T}) \cap B\right| < t\right] < 2e^{-\frac{l}{8|R_X|}}$$

Proof: Consider the indicator variables I_z taking a 1 when $A(z) = \mathsf{T}$, $\Pr_z[I_z = 1] = \frac{1}{|R_X|}$. Note that for $B \in \mathcal{Z} = \binom{Z}{l}$, $|A^{-1}(\mathsf{T}) \cap B| = \sum_{z \in B} I_z$, and the expectation of this sum is $l/|R_X|$. The above Chernoff bound directly (taking $p = \frac{1}{|R_X|}, \mu = \frac{l}{|R_X|}, \delta = \frac{1}{2}$) gives

$$\Pr_{z_1,..,z_l \in Z}\left[\sum_{i \in [l]} I_{z_i} < t = \frac{1}{2} \cdot l/|R_X|\right] < e^{-\frac{l}{8|R_X|}}$$

We are almost done, except that the above probability was taken with repetitions, while in our case, for $z_1, \ldots, z_l$ to constitute a block $B \in \mathcal{Z}$, they must be l distinct values. In fact, this happens with overwhelming probability and in particular $\geq \frac{1}{2}$, thus we write,

$$\begin{aligned}\Pr_{z_1,...,z_l \in Z}\left[\sum_i I_{z_i} < t \;\middle|\; |\{z_1,...,z_l\}| = l\right] &\leq \frac{\Pr_{z_1,...,z_l \in Z}\left[\sum_i I_{z_i} < t\right]}{\Pr_{z_1,...,z_l \in Z}\left[|\{z_1,...,z_l\}| = l\right]} \\ &\leq \frac{e^{-\frac{l}{8|R_X|}}}{\frac{1}{2}} = 2e^{-\frac{l}{8|R_X|}}\end{aligned}$$

■

Bibliography

[ABSS93] S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes and linear equations. In *Proc. 34th IEEE Symp. on Foundations of Computer Science*, pages 724–733, 1993.

[AD97] Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 284–293, El Paso, Texas, 4–6 May 1997.

[Ajt96] M. Ajtai. Generating hard instances of lattice problems. In *Proc. 28th ACM Symp. on Theory of Computing*, pages 99–108, 1996.

[Ajt98] Miklós Ajtai. The shortest vector problem in *L2* is NP-hard for randomized reductions. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC-98)*, pages 10–19, New York, May 23–26 1998. ACM Press.

[AK97] Rudolf Ahlswede and Levon H. Khachatrian. The complete intersection theorem for systems of finite sets. *European J. Combin.*, 18(2):125–136, 1997.

[ALM$^+$92] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and intractability of approximation problems. In *Proc. 33rd IEEE Symp. on Foundations of Computer Science*, pages 13–22, 1992.

[AS92] S. Arora and S. Safra. Probabilistic checking of proofs: A new characterization of NP. In *Proc. 33rd IEEE Symp. on Foundations of Computer Science*, pages 2–13, 1992.

[Bab86] L. Babai. On Lovász's lattice reduction and the nearest lattice point problem. *Combinatorica*, 6:1–14, 1986.

[BGLR93] M. Bellare, S. Goldwasser, C. Lund, and A. Russell. Efficient multi-prover interactive proofs with applications to approximation problems. In *Proc. 25th ACM Symp. on Theory of Computing*, pages 113–131, 1993.

[BGS98] Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, PCPs, and nonapproximability—towards tight results. *SIAM Journal on Computing*, 27(3):804–915, June 1998.

[BK97] J. Bourgain and G. Kalai. Influences of variables and threshold intervals under group symmetries. *Geom. Funct. Anal.*, 7(3):438–461, 1997.

[BKS99] Itai Benjamini, Gil Kalai, and Oded Schramm. Noise sensitivity of Boolean functions and applications to percolation. *Inst. Hautes Études Sci. Publ. Math.*, (90):5–43 (2001), 1999.

[Bol86] Béla Bollobás. *Combinatorics*. Cambridge University Press, Cambridge, 1986. Set systems, hypergraphs, families of vectors and combinatorial probability.

[BOL89] Ben-Or and Linial. Collective coin flipping. *ADVCR: Advances in Computing Research*, 5, 1989.

[BOLS88] M. Ben-Or, N. Linial, and M. Saks. Collective coin flipping and other models of imperfect randomness. In *Combinatorics (Eger, 1987)*, pages 75–112. North-Holland, Amsterdam, 1988.

[BYE85] R. Bar-Yehuda and S. Even. A local-ratio theorem for approximating the weighted vertex cover problem. *Annals of Discrete Mathematics*, 25:27–45, 1985.

[CN98] J.Y. Cai and A. Nerurkar. Approximating the SVP to within a factor $(1 + 1/\dim^\epsilon)$ is NP-hard under randomized reductions. In *Proc. of the 13th Annual IEEE Conference on Computational Complexity*, pages 46–55. 1998.

[Coo71] S. Cook. The complexity of theorem-proving procedures. In *Proc. 3rd ACM Symp. on Theory of Computing*, pages 151–158, 1971.

[DFK$^+$99] I. Dinur, E. Fischer, G. Kindler, R. Raz, and S. Safra. PCP characterizations of NP: Towards a polynomially-small error-probability. In *Proc. 31st ACM Symp. on Theory of Computing*, 1999.

[DKRS99] I. Dinur, G. Kindler, R. Raz, and S. Safra. Approximating-CVP to within almost-polynomial factors is NP-hard. Manuscript, 1999.

[DKS98] I. Dinur, G. Kindler, and S. Safra. Approximating-CVP to within almost-polynomial factors is NP-hard. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*, 1998.

[DS01] I. Dinur and S. Safra. The importance of being biased. In preparation, 2001.

[EH00] Lars Engebretsen and Jonas Holmerin. Clique is hard to approximate within $n^{1-o(1)}$. In *Automata, languages and programming (Geneva, 2000)*, pages 2–12. Springer, Berlin, 2000.

[EKR61] P. Erdős, Chao Ko, and R. Rado. Intersection theorems for systems of finite sets. *Quart. J. Math. Oxford Ser. (2)*, 12:313–320, 1961.

[ER60] P. Erdős and R. Rado. Intersection theorems for systems of sets. *J. London Math. Soc.*, 35:85–90, 1960.

[FF91] Peter Frankl and Zoltán Füredi. Beyond the Erdős-Ko-Rado theorem. *J. Combin. Theory Ser. A*, 56(2):182–194, 1991.

[FGL$^+$91] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. Approximating clique is almost NP-complete. In *Proc. 32nd IEEE Symp. on Foundations of Computer Science*, pages 2–12, 1991.

[FK94] U. Feige and J. Kilian. Two prover protocols–low error at affordable rates. In *Proc. 26th ACM Symp. on Theory of Computing*, pages 172–183, 1994.

[FK96] Ehud Friedgut and Gil Kalai. Every monotone graph property has a sharp threshold. *Proc. Amer. Math. Soc.*, 124(10):2993–3002, 1996.

[Fra78] P. Frankl. The Erdős-Ko-Rado theorem is true for $n = ckt$. In *Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976), Vol. I*, pages 365–375. North-Holland, Amsterdam, 1978.

[Fri98] Ehud Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, 18(1):27–35, 1998.

[GG98] O. Goldreich and S. Goldwasser. On the limits of non-approximability of lattice problems. In *Proc. 30th ACM Symp. on Theory of Computing*, pages 1–9, 1998.

[GMSS99] O. Goldreich, D. Micciancio, S. Safra, and J.-P. Seifert. Approximating shortest lattice vectors is not harder than approximating closest lattice vectors. *Inform. Process. Lett.*, 71(2):55–61, 1999.

[Hal00] Eran Halperin. Improved approximation algorithms for the vertex cover problem in graphs and hypergraphs. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 329–337, N.Y., January 9–11 2000. ACM Press.

[Hås97] Johan Håstad. Some optimal inapproximability results. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 1–10, El Paso, Texas, 4–6 May 1997.

[Hås99] Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Math.*, 182(1):105–142, 1999.

[Kar72] R. M. Karp. Reducibility among combinatorial problems. In Miller and Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.

[Kat68] G. Katona. A theorem of finite sets. In *Theory of graphs (Proc. Colloq., Tihany, 1966)*, pages 187–207. Academic Press, New York, 1968.

[Kho01] S. Khot. Improved inapproximability results for Max-Clique, Chromatic Number and Approximate Graph Coloring. In *Proc. 42nd IEEE Symp. on Foundations of Computer Science*, 2001.

[KKL88] J. Kahn, G. Kalai, and N. Linial. The influence of variables on Boolean functions. In IEEE, editor, *29th annual Symposium on Foundations of Computer Science, October 24–26, 1988, White Plains, New York*, pages 68–80. IEEE Computer Society Press, 1988.

[KLS00] Sanjeev Khanna, Nathan Linial, and Shmuel Safra. On the hardness of approximating the chromatic number. *Combinatorica*, 20(3):393–415, 2000.

[KMS98] David Karger, Rajeev Motwani, and Madhu Sudan. Approximate graph coloring by semidefinite programming. *J. ACM*, 45(2):246–265, 1998.

[Kru63] Joseph B. Kruskal. The number of simplices in a complex. In *Mathematical optimization techniques*, pages 251–278. Univ. of California Press, Berkeley, Calif., 1963.

[Lev73] L. Levin. Universal'nyĭe perebornyĭe zadachi (universal search problems : in Russian). *Problemy Peredachi Informatsii*, 9(3):265–266, 1973.

[LLL82] A.K. Lenstra, H.W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261:513–534, 1982.

[LLS90] J. Lagarias, H.W. Lenstra, and C.P. Schnorr. Korkine-Zolotarev bases and successive minima of a lattice and its reciprocal lattice. *Combinatorica*, 10:333–348, 1990.

[LY94] Carsten Lund and Mihalis Yannakakis. On the hardness of approximating minimization problems. *Journal of the ACM*, 41(5):960–981, 1994.

[Mar74] G. A. Margulis. Probabilistic characteristics of graphs with large connectivity. *Problemy Peredači Informacii*, 10(2):101–108, 1974.

[Mic98] D. Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. In *Proc. 39th IEEE Symp. on Foundations of Computer Science*, 1998.

[MR97] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, 1997.

[MS83] B. Monien and E. Speckenmeyer. Some further approximation algorithms for the vertex cover problem. In G. Ausiello and M. Protasi, editors, *Proceedings of the 8th Colloquium on Trees in Algebra and Programming (CAAP'83)*, volume 159 of *LNCS*, pages 341–349, L'Aquila, Italy, March 1983. Springer.

[PY91] C. Papadimitriou and M. Yannakakis. Optimization, approximation and complexity classes. *Journal of Computer and System Sciences*, 43:425–440, 1991.

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, June 1998.

[RS97] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proc. 29th ACM Symp. on Theory of Computing*, pages 475–484, 1997.

[Rus82] Lucio Russo. An approximate zero-one law. *Z. Wahrsch. Verw. Gebiete*, 61(1):129–139, 1982.

[Sch85] C.P. Schnorr. A hierarchy of polynomial-time basis reduction algorithms. In *Proceedings of Conference on Algorithms, Pécs (Hungary)*, pages 375–386. North-Holland, 1985.

[vEB81] P. van Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. Technical Report 81–04, Math. Inst. Univ. Amsterdam, 1981.

[Wil84] Richard M. Wilson. The exact bound in the Erdős-Ko-Rado theorem. *Combinatorica*, 4(2-3):247–257, 1984.