ON THE SECURITY OF GOLDREICH'S ONE-WAY FUNCTION

Andrej Bogdanov and Youming Qiao

Abstract. Goldreich (ECCC 2000) suggested a simple construction of a candidate one-way function $f : \{0,1\}^n \to \{0,1\}^m$ where each bit of output is a fixed predicate P of a constant number d of (random) input bits. We investigate the security of this construction in the regime $m = Dn$, where $D(d)$ is a sufficiently large constant. We prove that for any predicate P that correlates with either one or two of its inputs, f can be inverted with high probability.
We also prove an amplification claim regarding Goldreich's construction. Suppose we are given an assignment $x' \in \{0,1\}^n$ that has correlation $\varepsilon > 0$ with the hidden assignment $x \in \{0,1\}^n$. Then, given access to x', it is possible to invert f on x with high probability, provided $D = D(d, \varepsilon)$ is sufficiently large.

Keywords. one-way function, parallel cryptography

Subject classification. 68Q17

1. Introduction

Oded Goldreich (Goldreich 2000a) proposed a very simple construction of a conjectured one-way function:

1. From the family of bipartite graphs with n vertices on the left, m vertices on the right, and regular right-degree d, randomly choose a graph G.

2. From all predicates mapping $\{0,1\}^d$ to $\{0,1\}$, randomly choose a predicate P.

3. Based on the chosen graph G and predicate P, let $f = f_{G,P}$ be the function from $\{0,1\}^n$ to $\{0,1\}^m$ defined by

$$f(x)_j = \text{the } j\text{th bit of } f(x) = P(x_{\Gamma(j,1)}, \ldots, x_{\Gamma(j,d)})$$

where $\Gamma(j,k)$ is the kth neighbor of right vertex j of G.

Goldreich conjectured that when $m = n$ and d is constant, for most graphs G and predicates P, the resulting function is one-way.[1]

In this work we investigate Goldreich's construction in the setting where the graph G is random, d is constant, and $m = Dn$ for a sufficiently large constant $D = D(d)$. We show that for this setting of parameters, Goldreich's construction is not secure for predicates correlating with one input or with a pair of the inputs. (As d increases, most predicates are of this type.)

We also show that if we are given a "hint" x' – any assignment that has nontrivial correlation with the actual input x to the one-way function – it is possible to invert f on x, as long as D is a sufficiently large constant which depends on both d and the correlation between x and x'.

Our results indicate that the security of Goldreich's construction is fairly sensitive on the output to input length ratio m/n. We show that when m/n is a sufficiently large constant (depending on d), for a large class of predicates the function can be inverted on a large fraction of inputs. It is also known that when m/n is smaller than $1/(d-1)$ the function can be inverted for every predicate P, since with high probability the "constraint hypergraph" splits into components of size $O(\log n)$ (Schmidt & Shamir 1985).

Our analysis leaves open the possibility that for specific choices of P that fall outside our characterization, the function is one-way even when the output is much longer than the input. Consider any predicate P which is balanced, does not correlate with any of its inputs, and does not correlate with any pair of its inputs. Even when m is substantially larger than n, say $m = n^{1.1}$, we do not know of any method for inverting Goldreich's function based on the predicate P. In fact, we do not even know whether the output of this function can be efficiently distinguished from a random string of length m.

1.1. Goldreich's function and Cryptography in NC^0. Goldreich's proposal for a one-way function has several features that were absent from earlier proposals: (1) It is extremely simple to implement, and (2) it is very fast to compute, especially in parallel. On the other hand, the conjectured security of Goldreich's function is not known to relate to any standard assumptions in cryptography, such as hardness of factoring or hardness of finding short vectors in lattices.

The design of cryptographic constructions in NC^0 (i.e., constructions where every output bit depends on a constant number of input bits) has since been

[1] More precisely, Goldreich conjectures that for any fixed family of graphs $\{G_n\}$ with certain expansion properties and most predicates P on d bits, the family of functions $\{f_{G_n,P}\}$ is one-way.

extended to other cryptographic primitives, in particular pseudorandom generators. Remarkably, Applebaum *et al.* (2004) showed that pseudorandom generators (and in particular one-way functions) in NC^0 can be obtained under many commonly used assumptions (such as hardness of discrete logarithm, hardness of factoring, and hardness of finding short vectors in lattices). In a different work, Applebaum *et al.* (2006) gave a different construction of a pseudorandom generator with linear stretch using the less standard assumption that certain random linear codes are hard to decode.

The results of Applebaum *et al.* (2004, 2006) are obtained by starting with known constructions of cryptographic primitives that reside outside NC^0, and transforming them into NC^0 variants that are secure under the same hardness assumption. These transformations entail a loss of parameters. To yield reasonable hardness, these constructions require fairly large input length. Also, pseudorandom generators obtained using this process have only small linear stretch. It is not known whether a pseudorandom generator that stretches n bits of input into, say, $n^{1.1}$ bits of output can be obtained under similar assumptions.

For this reason, we believe that it is interesting to investigate direct constructions of pseudorandom primitives in NC^0, which have the potential to yield better parameters. In this direction, Mossel *et al.* (2003) proposed the construction of a pseudorandom generator in NC^0 with potentially superlinear stretch. They proved that for any constant c, there is a function in NC^0 that maps n bits to n^c bits and is pseudorandom against all linear tests.

More recently, Applebaum *et al.* (2010) showed that for certain choices of the predicate P, Goldreich's function (with slightly superlinear stretch) is pseudorandom against linear tests, low-degree polynomial tests, and tests implemented by polynomial-size constant-depth circuits.

Cook *et al.* (2009) showed that a restricted class of algorithms called "myopic algorithms" take exponential time to invert Goldreich's construction. The kinds of algorithms used in this work are not myopic.[2]

1.2. Our Results. We state our main results. We say that algorithm A *inverts* function $f\colon \{0,1\}^n \to \{0,1\}^m$ on input x if $A(f, f(x))$ returns a preimage for $f(x)$, where $x \in \{0,1\}^n$. For our second application we will also allow A to take as part of its input some auxiliary information. Our definition is slightly non-standard (see Definition 2.4.3 in Goldreich (2000b)) as the inverter takes the description of the function it is supposed to invert as part of its input.

[2]Out of the algorithms used here, only the algorithm from Section 3.1 can be naturally viewed as myopic.

This is convenient when working with Goldreich's function on a random graph (which is essentially a non-uniform function family), as it allows for the description of the graph to be furnished to the algorithm.

To state the theorems, we need to define three standard combinatorial properties of a predicate $P\colon \{0,1\}^d \to \{0,1\}$. The *single variable correlation* of P is the quantity

$$\gamma_1(P) = \max_{i\in[d]} |\Pr[P(z) = z_i] - \Pr[P(z) \neq z_i]|.$$

The *pairwise correlation* of P is the quantity

$$\gamma_2(P) = \max_{i\neq j, i,j\in[d]} |\Pr[P(z) = z_i \oplus z_j] - \Pr[P(z) \neq z_i \oplus z_j]|.$$

The *boundary* of P is the quantity

$$\beta(P) = \Pr_{z\sim\{0,1\}^d}[\exists z', |z - z'| = 1\colon P(z) \neq P(z')].$$

In all cases, z is chosen uniformly at random from $\{0,1\}^d$. Notice that the values of these quantities are multiples of 2^{-d}. Moreover, for any non-constant predicate, $\beta(P)$ is nonzero. It is well known that if P is balanced then $\beta(P) = \Omega(d^{-1/2})$.

We now state our main theorems. The first two theorems give inversion algorithms for predicates that correlate with one or a pair of the inputs, provided that the output length to input length ratio is sufficiently large. All of the theorems refer to the function family $f_{G,P}\colon \{0,1\}^n \to \{0,1\}^m$ with $m = Dn$ and $P\colon \{0,1\}^d \to \{0,1\}$. The quantities γ_1, γ_2, and β refer to the predicate P.

THEOREM 1.1. *Let K be a sufficiently large constant. Assume that $\gamma_1 > 0$, $d \leq \beta n^{1/K}/K$, and $D \geq \max\{K/\gamma_1^2 \log(d/\beta), (d/\beta)^K\}$. Then for every $r \leq (\beta/d)^K n$, there exists an algorithm that runs in time $D^3 n^{O(r)}$ and inverts $f_{G,P}(x)$ for a $1 - O(n^{-r})$ fraction of pairs (G, x).*

Theorem 1.1 gives a family of algorithms that exhibit a tradeoff between running time and success probability. When r is a constant, the inverter runs in polynomial time and succeeds on an inverse polynomial fraction of inputs. Also, observe that while the output is required to be always longer than the input, when γ_1 and β are not too small, for instance inverse polynomial in d, then the inverter succeeds even when $m = \mathrm{poly}(d) \cdot n$.

THEOREM 1.2. *Let K be a sufficiently large constant. Assume that $\gamma_2 > 0$, $d \leq \beta n^{1/K}/K$, and $D \geq K(d/\beta\gamma_2)^K$. Then for every $r \leq (\beta/d)^K n$, there exists an algorithm that runs in time $D^3 n^{O(r)}$ and inverts $f_{G,P}(x)$ for a $1 - O(n^{-r})$ fraction of pairs (G, x).*

Our last theorem gives an inversion algorithm that applies to all predicates, but requires knowledge of a pre-image x' of $f_{G,P}(x)$ that correlates with x (the formal definition of correlation between a pair of assignments is given in Section 2). In this theorem, we require D to be at least polynomial in $(1/\varepsilon)^d$.

THEOREM 1.3. *Let K be a sufficiently large constant, $\varepsilon > 0$, and $D \geq 2^{Kd}/\varepsilon^{2d-2}$. Let $P : \{0,1\}^d \to \{0,1\}$ be any predicate. Then for every $r \leq 2^{-Kd} n$, there exists an algorithm that runs in time polynomial in D and n^r with the following property. For a $1 - O(n^{-r})$ fraction of pairs (G, x) and every assignment x' that has correlation at least ε (in absolute value) with x, on input $(f_{G,P}, f_{G,P}(x), x')$, A outputs an inverse for $f_{G,P}(x)$.*

Our results also generalize to the case where different predicates are used to compute different bits of the output. To simplify the presentation, we restrict our proofs to the case when the same predicate is used.

1.3. Our Approach. The problem of inverting Goldreich's function is somewhat analogous to the problem of reconstructing assignments to random 3SAT formulas in the planted 3SAT model. We exploit this analogy and show that several of the tools developed for planted 3SAT can be applied to our setting as well.

The proofs of our theorems are carried out in two stages. In the first stage, we almost invert f in the sense that we find an assignment z that matches the hidden assignment x on a 99% fraction of positions. In the second stage we turn z into a true inverse for $f(x)$. The second stage is common to the proofs of all theorems.

To give some intuition about the first stage in Theorem 1.1, suppose for instance that P is the majority predicate. Then we try to guess a the value of the input bit x_i by looking at all constraints where x_i appears and taking the majority of these values. Since x_i has positive correlation with the majority predicate, we expect this process to result in a good guess for most x_i that appear in a sufficiently large number of clauses. In fact, if f has about $n \log n$ bits of output, this reconstructs the assignment completely; if $m = Dn$ for a sufficiently large constant D, a large constant fraction of the bits of x is recovered. This argument, which applies to any predicate that correlates to one of its inputs, is given in Section 3.1.

For Theorem 1.2, we view each output of f as a noisy indicator for the value of the parity of the pair of inputs it correlates with. This allows us to write a noisy system of linear equations with two variables per equation whose intended solution is the hidden assignment x. Using an approximation algorithm for such systems (Charikar & Wirth 2004; Goemans & Williamson 1995), we can extract an assignment x' that correlates with x. We then show that the correlation between x and x' can be improved via a self-correction step, which takes advantage of certain expansion properties of the system of equations that follow (with high probability) from the randomness of G.[3]

The first stage in the proof of Theorem 1.3 is based on the observation that if we start with some assignment x' that correlates with the input x to f, then the output bits of $f(x)$ give information about the values of various inputs x_i, for an arbitrary predicate P. We prove this in Section 5. This correlation amplification procedure works in a more general setting than the one used in the proof of Theorem 1.2, but yields worse parameters.

For the second stage, we base our algorithm on known approaches for finding solutions of planted random instances. Alon & Kahale (1997) showed how to find a planted 3-coloring in a random graph of constant degree. Flaxman (2003) (see also Krivelevich & Vilenchik (2006); Vilenchik (2007)) gave a similar algorithm for finding an assignment in a planted random 3SAT formula with sufficiently large clause-to-variable ratio. The planted 3SAT model can be viewed as a variant of our model where the predicate P corresponds to one of the eight predicates $z_1 \vee z_2 \vee z_3, \ldots, \overline{z_1} \vee \overline{z_2} \vee \overline{z_3}$. This algorithm starts from an almost correct assignment, then unsets a small number of the variables in this assignment according to some condition ("small support size"), so that with high probability all (but a constant number of) the remaining set variables are correct. Then the value of the unset variables can be inferred in polynomial time. We show that the notion of "small support size" can be generalized to arbitrary non-constant predicates, and this type of algorithm can be used to invert f. While we directly follow previous approaches, our proofs in Section 4 include some technical simplifications and follow a more rigorous presentation style.

[3] This algorithm was suggested to us by Benny Applebaum. Our initial solution was based on a spectral partitioning algorithm for random graphs, but we chose to present the suggested solution owing to its technical simplicity and improved parameters.

2. Definitions and notation

Let X, Y be random variables over $\{0,1\}$. The *correlation* between X and Y is the value

$$\Pr[X = Y] - \Pr[X \neq Y] = 2(\Pr[X = Y] - 1/2).$$

The correlation between a predicate $P\colon \{0,1\}^d \to \{0,1\}$ and its ith input is the correlation between the random variables $P(X_1, \ldots, X_d)$ and X_i. The correlation between P and the pair of inputs i, j is the correlation between the random variables $P(X_1, \ldots, X_d)$ and $X_i \oplus X_j$. Here $X_1, \ldots, X_d$ are chosen uniformly at random. We say P correlates with its ith input (resp. with the pair of inputs i, j) if the above correlation is non-zero.

The correlation between a pair of assignments $x, y \in \{0,1\}^n$ is defined as the correlation between the ith bit of x and y, where i is chosen uniformly at random from $[n]$.

We say an assignment $x \in \{0,1\}^n$ is ε-balanced if $|\Pr[x_i = 0] - 1/2| \leq \varepsilon$. A Bernoulli random variable $X \sim \{0,1\}$ is ε-biased towards 0 (resp. 1) if $\Pr[X = 0]$ is no less than $1/2 + \varepsilon$ (resp. no more than $1/2 - \varepsilon$).

For two random variables X, Y over the same finite domain Ω their *statistical distance* $\mathrm{sd}(X, Y)$ is the quantity $\frac{1}{2}\sum_{\omega\in\Omega}(|X(\omega) - Y(\omega)|)$.

The random graph model In our random graph model, the bipartite graph G in the function $f_{G,P}$ is chosen from the following random graph model $\mathcal{G} = \{\mathcal{G}_{n,m,d}\}$: (1) Each graph G in $\mathcal{G}_n$ has n left vertices and $m = m(n)$ right vertices; (2) each right vertex v of G has d neighbors on the left, labeled by $\Gamma(v, 1), \ldots, \Gamma(v, d)$; (3) The neighbors of each right vertex are uniformly distributed (repetitions allowed) and independent of the neighbors of all other vertices.

One can also consider variants of the model where repeated neighbors are not allowed (as in Goldreich's original proposal (Goldreich 2000a)), or a where for each d-tuple of inputs the corresponding output is present independently with probability $p = p(n)$ (as is common in the planted SAT literature). Our results extend to such variants.

3. Obtaining an Almost Correct Assignment

In this section, we show that when the predicate P correlates with one or two of its inputs, it is possible (with high probability) to approximately invert $f_{G,P}(x)$, namely find an assignment x' that agrees with x on almost all inputs.

3.1. Predicates Correlating with One Input. When the predicate P correlates with one of its inputs, then every output bit of $f_{G,P}(x)$ gives an indication about what the corresponding input bit should be. If we think of this indication as a vote, and take a majority of all the votes, we set most of the input bits correctly. The following algorithm and proposition formalize this idea.

Recall that $\gamma_1(P)$ denotes the maximum correlation (in absolute value) between P and one of its inputs. Without loss of generality, we will assume that this correlation is attained by z_1, and that the correlation is positive. (If the correlation is negative, we can work with function obtained by complementing each output of $f_{G,P}$.)

Algorithm Single Variable Correlation:
Input: A predicate P; a graph G; the value $f_{G,P}(x) \in \{0,1\}^m$.

1. Let $\nu = \Pr[P(z) = 1]$. For every input i, set x'_i to 1 if at least a ν fraction of the values $f_{G,P}(x)_j$ where i occurs as the first input in $f_{G,P}(x)_j$ evaluate to 1, and 0 otherwise.

2. Output the assignment x'.

Proposition 3.1. *Assume the correlation $\gamma_1(P)$ is attained between P and its first input and this correlation is positive. Assume also that $\varepsilon > 0$ and $D > (16/\gamma_1^2)\log(4/\varepsilon)$. For every x that is $\gamma_1/4d$-balanced, with probability $1 - 2^{-\Omega(\varepsilon n)}$ over the choice of G, on input $f(x)$, the assignment x' produced by algorithm* Single Variable Correlation *agrees with x on a $(1-\varepsilon)$ fraction of inputs.*

By the Chernoff bound, all but a $2^{-\Omega(\gamma_1^2 n/d^2)}$ fraction of inputs $x \in \{0,1\}^n$ are $\gamma_1/4d$-balanced.

To explain the proof, let us make the unrealistic assumption that x is perfectly balanced. When the graph G is random, on average the i-th bit of x will be represented in dD of the outputs. Out of those, on average it will figure in the first position in D outputs. From the perspective of each of these outputs, the other input bits are chosen uniformly at random (because x is balanced), and we expect each one of them to exhibit some correlation towards the input. The amount of correlation has to be at least $\gamma_1(P)$, so if D is sufficiently large, on average the effect of x_i on the outputs where it is involved as a first variable becomes noticeable and the outputs can be used to predict the value of x_i with good probability.

To make the argument precise, we must argue that this average-case behavior is representative for most of the input bits. To do so we use an application of the Chernoff bound which is tailored to our setting and is given in Lemma A.1 in Appendix A. We also need to deal with the unrealistic assumption that x is perfectly balanced; to do so, we replace it by the weaker assumption that x is almost balanced, which is satisfied by most $x \in \{0,1\}^n$.

PROOF. Fix an x that is $\gamma_1/2d$-balanced. Let N_i be the number of constraints whose first input is in i, and I be the set of those inputs i for which $N_i \geq D/2$. We first show that conditioned on the numbers N_i, with probability at least $1 - 2^{-\Omega(\varepsilon n)}$, $x'_i = x_i$ for all but $\varepsilon n/2$ of the inputs $i \in I$. We then show that with the same probability, I has size at least $(1 - \varepsilon/2)n$.

Let us fix $i \in I$ and upper bound the probability that $x'_i \neq x_i$. Every output that involves i as its first variable is sampled from the distribution $P(x_i, \tilde{z}_2, \ldots, \tilde{z}_d)$, where $\tilde{z}_2, \ldots, \tilde{z}_d$ are independent bernoulli random variables that are $\gamma_1/4d$-biased. Replacing each $\tilde{z}_k$ by a uniformly random bit $z_k \sim \{0,1\}$ affects the distribution of $P(\cdot)$ by at most $\gamma_1/2d$, so by the triangle inequality

$$|\Pr[P(x_i, \tilde{z}_2, \ldots, \tilde{z}_d) = x_i] - \Pr[P(x_i, z_2, \ldots, z_d) = x_i]| \leq \frac{\gamma_1}{4}.$$

From the definitions of ν and γ_1 we get that

$$\Pr[P(1, z_2, \ldots, z_d) = 1] = \nu + \frac{\gamma_1}{2} \quad \text{and} \quad \Pr[P(0, z_2, \ldots, z_d) = 0] = \nu - \frac{\gamma_1}{2}$$

and so it follows that

$$\Pr[P(1, \tilde{z}_2, \ldots, \tilde{z}_d) = 1] = \nu + \frac{\gamma_1}{4} \quad \text{and} \quad \Pr[P(0, \tilde{z}_2, \ldots, \tilde{z}_d) = 0] = \nu - \frac{\gamma_1}{4}.$$

By the Chernoff bound, the probability that $x'_i \neq x_i$ is at most $2^{-\gamma_1^2 N_i/8} \leq 2^{-\gamma_1^2 D/16} \leq \varepsilon/4$.

Conditioned on the numbers $N_i, i \in I$, the events $x'_i \neq x_i$ are independent of one another, because once the first input of every output is revealed, the other inputs that participate in the prediction of x'_i are chosen independently of one another. Since for each $i \in I$, the probability of the event $x'_i \neq x_i$ is at most $\varepsilon/4$, by the Chernoff bound the number of inputs $i \in I$ such that $x'_i \neq x_i$ is at most $\varepsilon/2$ with probability at least $1 - 2^{-\Omega(\varepsilon n)}$.

By Lemma A.1 in Appendix A, the probability that I has size less than $(1 - \varepsilon/2)n$ is at most $2^{-\Omega(\varepsilon D n)}$. So with probability $1 - 2^{-\Omega(\varepsilon D n)}$, $x'_i = x_i$ for all but at most $\varepsilon n/2$ inputs inside I and $\varepsilon n/2$ inputs outside I. The proposition follows. □

3.2. Predicates Correlating with a Pair of Inputs. Let us now assume that the predicate P correlates with a pair of its inputs. Without loss of generality, we may assume that P correlates with the first pair of inputs (z_1, z_2), and that the correlation is positive (otherwise, we can work with the complement of P by complementing all the outputs of $f_{G,P}(x)$):

$$\Pr[P(z) = z_1 \oplus z_2] \geq \frac{1}{2} + \frac{\gamma_2}{2}.$$

We can then think of each output $f(x)_j$ as giving noisy information about the value $x_{\Gamma(j,1)} \oplus x_{\Gamma(j,2)}$. If x is balanced, on average a $1/2 + \gamma_2/2$ fraction of the linear equations $x_{\Gamma(j,1)} \oplus x_{\Gamma(j,2)} = f(x)_j$ will be satisfied. If x is almost balanced, we still expect $1/2 + \Omega(\gamma_2)$ of them to be satisfied. Using an approximation algorithm of Charikar and Wirth, we can obtain a solution x' that satisfies $1/2 + \Omega(\gamma_2/\log(1/\gamma_2))$ fraction of these equations:

THEOREM 3.2 (Charikar & Wirth 2004). *There is a randomized algorithm CW that given a system of m linear equations modulo 2 and a parameter $\delta > 0$, finds an assignment that satisfies at least $m/2 + \Omega(\delta m/\log(1/\delta))$ of the equations, provided that $m/2 + \delta m$ of the equations can be satisfied simultaneously. The algorithm runs in expected time polynomial in m/δ.*

We will argue that (with high probability over the choice of G) (1) x' must correlate with x and (2) This correlation can be amplified significantly by applying a round of self-correction to this system of equations. Specifically, we show that with high probability over G, the following algorithm recovers most of the bits of x:

ALGORITHM Pairwise Correlation:
Input: A predicate P; a graph G; the value $f_{G,P}(x) \in \{0,1\}^m$.

1. Create the following system of equations: For every $j \in [m]$,

   $$u_{\Gamma(j,1)} \oplus u_{\Gamma(j,2)} = f(x)_j \tag{3.3}$$

   Let H be a directed graph over vertex set $[n]$ with m edges $(\Gamma(j,1), \Gamma(j,2))$.

2. Apply algorithm CW on input (3.3) and $\gamma_2/8$ to obtain an assignment $x' \in \{0,1\}^n$.

3. (Self-correction) For every $i_1 \in [n]$, calculate the number Q_{i_1} of equations $(i_1, i_2) = (\Gamma(j,1), \Gamma(j,2))$ where $f(x)_j = x'_{i_2}$. Sort the variables by order

of increasing Q_i, breaking ties arbitrarily. Output the assignments $y^{(k)}$ and $\overline{y^{(k)}}$, $k \in \{0, \ldots, n\}$ where

$$y_i^{(k)} = \begin{cases} 1, & \text{if } i \text{ is among the } k \text{ variables with smallest value } Q_i, \\ 0, & \text{otherwise.} \end{cases}$$

and $\overline{y^{(k)}}$ is the complementary assignment obtained by swapping 0 and 1.

In step 3, it may be helpful to think of the value $x'_{i_2} \oplus f(x)_j$ in the equation $u_{i_1} \oplus u_{i_2} = f(x)_j$ as a "vote" that x_{i_1} should take value zero. The quantity Q_{i_1} tallies the votes for x_{i_1} from all the equations that involve i_1 as the first input. The next step would be to set a threshold t so that all inputs with $Q_i > t$ are set to zero, and the others are set to one. A natural threshold to use is the median value of $Q_1, \ldots, Q_n$. While this would be sufficient to prove correctness, it would require us to make somewhat stronger assumptions about the balance of x and would introduce technical complications in the analysis. Instead, we consider every possible threshold, which produces $n + 1$ candidate assignments $y^{(0)}, \ldots, y^{(n)}$. One additional complication is that the correlation effects may be negative, in which case Q_{i_1} should be interpreted as a vote for $x_{i_1} = 1$ and not $x_{i_1} = 0$. This suggests that we should also consider the negated assignments $\overline{y^{(k)}}$ as possible solutions.

Abusing terminology we use "edge j" to refer to the edge $(\Gamma(j, 1), \Gamma(j, 2))$.

Proposition 3.4. *Assume the correlation $\gamma_2(P)$ is attained between P and its first pair of inputs and this correlation is positive. Assume also that $\gamma_2(P) > K(\gamma_1(P))^{2/3}$ and $D \geq K(\log(1/\gamma_2)/\gamma_2)^2$ for a sufficiently large constant K. For every x that is $\gamma_2^{3/2}/(12\log(1/\gamma_2)d)$-balanced and with probability $1 - 2^{-n}$ over the choice of G, on input $f(x)$, at least one of the assignments produced by algorithm* Pairwise Correlation *agrees with x on all but a $O(\sqrt{\log(1/\gamma_2)}/(\sqrt{D}\gamma_2^{3/2}))$ fraction of inputs.*

By the Chernoff bound, all but $2^{-\Omega(\tilde{\gamma}_2^3 n/d^2)}$ inputs $x \in \{0,1\}^n$ are properly balanced, where $\tilde{\gamma_2} = \gamma_2/(\log(1/\gamma_2))^{2/3}$.

We outline the proof of Proposition 3.4. From the randomness of G it follows that with high probability, x satisfies $1/2 + \Omega(\gamma_2)$ of the equations (3.3), in which case x' will satisfy $1/2 + \tilde{\Omega}(\gamma_2)$ of the same equations. We will argue that x and x' must then have correlation $\tilde{\Omega}(\gamma_2)$. To see this, notice that the assignments x and x' differ in satisfying the equation $u_{i_1} \oplus u_{i_2} = f(x)_j$ exactly when $x_{i_1} \oplus x'_{i_1} \neq x_{i_2} \oplus x'_{i_2}$. If we think of the equation as an edge in H, then the differences are caused by those edges that cross the cut between

those inputs that take the same value in x and x' and those that take different values. With high probability, the graph H is expanding; so if the cut was almost balanced, about half of the edges would be crossing it.

Let us now make the unrealistic assumption that x satisfies all the equations. If the correlation between x and x' was small, then the cut would be almost balanced, so x' could satisfy only about half the equations. Since x' satisfies noticeably more than half the equations, it would follow that x and x' have noticeable correlation. However we merely know that x satisfies $1/2 + \Omega(\gamma_2)$ of the equations. It could then possibly happen that although the cut is balanced, most of the edges in the cut come from equations that are unsatisfied by x, in which case x' could end up satisfying substantially more than half the equations.

To show this is not possible, we would like to partition the edges of H into subgraphs H^1 and H^0, consisting of those edges induced by the equations satisfied and unsatisfied by x, respectively. Unfortunately, H^1 and H^0 may not be expanding. (For instance, if $P(z_1, z_2, z_3)$ is the predicate that is true if and only if $z_1 = z_2 = z_3$, the graph H^0 has an almost-balanced cut with no edges crossing it.) However, if we now partition the vertex set into $S_0 = \{i \colon x_i = 0\}$ and $S_1 = \{i \colon x_i = 1\}$, the restrictions of H^0 and H^1 on each of the cuts (S_{a_1}, S_{a_2}) will be random and therefore likely to be expanding. By applying the analysis to each of these subgraphs, we can still conclude that x and x' must be correlated.

At this point, it remains to amplify the correlation between x and x'. One possibility is to apply the generic amplification procedure from Section 5. However, we can obtain an improved analysis (specifically, a better dependence $D(d)$) for the special class of predicates that correlate with two variables. Let us look at a random equation $u_{i_1} \oplus u_{i_2} = f(x)_j$. On average, $f(x)_j$ is correlated with $x_{i_1} \oplus x_{i_2}$ and x'_{i_2} is also correlated with x_{i_2}, and since the two "noises" are independent, we would expect that x_{i_1} should be correlated with $f(x)_j \oplus x'_{i_2}$. By large deviation, we could then hope that the average value of $f(x)_j \oplus x'_{i_2}$ over all equations involving x_{i_1} should give significant information about x_{i_1}, allowing us to amplify the correlation. Thanks to the expansion of the graphs involved, we can show that this average behavior is typical for most of the inputs, allowing us to amplify the correlation between x and x' significantly.

We now introduce some notation. Partition the edges of H into subgraphs $H^b_{a_1 a_2}$, $a_1, a_2, b \in \{0, 1\}$ as follows. For each edge j of H where $x_{\Gamma(j,1)} = a_1$ and $x_{\Gamma(j,2)} = a_2$:

- $H^1_{a_1 a_2}$ contains edge j if $x_{\Gamma(j,1)} \oplus x_{\Gamma(j,2)} = f(x)_j$, and
- $H^0_{a_1 a_2}$ contains edge j if $x_{\Gamma(j,1)} \oplus x_{\Gamma(j,2)} \neq f(x)_j$.

Every edge from H is present in $H^b_{a_1a_2}$ with some probability $p^b_{a_1a_2}$. We begin by showing that with high probability all of the graphs $H^b_{a_1a_2}$ are expanding, and argue that in such a case x and x' must be correlated. We first give a general random graph model H^* that describes all of the graphs $H^b_{a_1a_2}$. Let H^* be a graph on vertex set $S^* \cup T^*$ (where $S^*, T^* \subseteq [n]$) chosen from the following distribution. For each of m possible edges, with probability p^*, choose random vertices $i_1 \in S^*, i_2 \in T^*$ and add the edge (i_1, i_2) to H^*. (Otherwise, do nothing.)

We will say H^* is η-*expanding* if for every pair of subsets $S \subseteq S^*, T \subseteq T^*$

$$\left| |\{\text{edges } (i_1, i_2) \text{ in } H^* : i_1 \in S, i_2 \in T\}| - \frac{|S|}{|S^*|}\frac{|T|}{|T^*|}p^*m \right| \leq \eta m.$$

Claim 3.5. *With probability* $1 - 2^{-2n}$, H^* *is* $2/\sqrt{D}$*-expanding.*

Proof. The expected number of edges from S to T is $\frac{|S|}{|S^*|}\frac{|T|}{|T^*|}p^*Dn$. Moreover, the events that each one of the Dn potential edges satisfies this property are independent, so by a Chernoff Bound, for a specific pair (S, T), the expression under the probability is at most $2/\sqrt{D}$ with probability at most $2e^{-4n} < 2^{-3n}$. The claim follows by taking a union bound over all pairs of sets (S, T). □

Claim 3.6. *For every* $x \in \{0,1\}^n$ *with probability* $1 - 2^{-2n}$, *the number of edges in* $H^b_{a_1a_2}$ *is within* $2\sqrt{D}n$ *of* $p^b_{a_1a_2}m$.

This claim follows from the Chernoff bound. Let $S_0 = \{i \colon x_i = 0\}$ and $S_1 = \{i \colon x_i = 1\}$. For the following two claims, we introduce a value $\alpha \in [-1, 1]$ that measures the correlation between x and x' defined as follows: First let $\alpha_0, \alpha_1 \in [-1, 1]$ be values that satisfy $\Pr[x'_i = a \mid x_i = a] = \frac{1}{2}(1+\alpha_a)$, $a \in \{0, 1\}$, and let $\alpha = \frac{1}{2}(\alpha_0 + \alpha_1)$. The following Claim relates the number of equations satisfied by x' to this correlation measure α.

Claim 3.7. *Assume* $H^b_{a_1a_2}$ *is* η*-expanding and has* $p^b_{a_1a_2}m \pm \eta m$ *edges for all* $a_1, a_2, b \in \{0, 1\}$. *Suppose* x' *satisfies* $\frac{1}{2}(1 + \gamma')m$ *of the equations (3.3). Then* $\gamma' \leq 2\alpha^2 + 24\eta$.

Proof. Suppose x' satisfies $\frac{1}{2}(1 + \gamma')m$ of the equations. Let $Z = \{i \colon x_i = x'_i\}$. Notice that x and x' differ in satisfying the jth equation if and only if

edge j crosses the cut $(Z, \overline{Z})$. By our expansion assumption, for every a_1, a_2, b we have

$$\left|\left|\{\text{edges } (i_1, i_2) \text{ in } H^b_{a_1 a_2} : i_1 \in Z, i_2 \in \overline{Z}\}\right| - \tfrac{1}{2}(1+\alpha_{a_1}) \cdot \tfrac{1}{2}(1-\alpha_{a_2}) \cdot p^b_{a_1 a_2} m\right| \leq \eta m.$$

since the density of Z in S_{a_1} is $\frac{1}{2}(1+\alpha_{a_1})$, and density of $\overline{Z}$ in S_{a_2} is $\frac{1}{2}(1+\alpha_{a_2})$. It follows that the number of equations satisfied by x' is at most

$$\begin{aligned}\tfrac{1}{2}(1+\gamma')m \leq & \sum_{a_1, a_2 \in \{0,1\}} (p^1_{a_1 a_2} m + \eta m) \\ & + 2 \sum_{a_1, a_2 \in \{0,1\}} \big(\tfrac{1}{2}(1+\alpha_{a_1}) \cdot \tfrac{1}{2}(1-\alpha_{a_2}) \cdot p^0_{a_1 a_2} m + \eta m\big) \\ & - 2 \sum_{a_1, a_2 \in \{0,1\}} \big(\tfrac{1}{2}(1+\alpha_{a_1}) \cdot \tfrac{1}{2}(1-\alpha_{a_2}) \cdot p^1_{a_1 a_2} m - \eta m\big).\end{aligned}$$

The first summation accounts for all the equations satisfied by x, while the other two account for those equations satisfied by x' but not x and those equations satisfied by x but not x' (with $x_{i_1} \in S_{a_1}$ and $x_{i_2} \in S_{a_2}$), respectively. The conclusion follows after simplifying this expression. □

Claim 3.8. *Assume $\alpha > 0$, $\alpha\gamma_2 > \gamma_1$, x is $\alpha\gamma_2/12d$-balanced, and H is η-expanding. Then there exists some $k \in [n]$ so that for all but a $O(\eta/\alpha\gamma_2)$ fraction of the inputs $i \in [n]$, $y_i^{(k)} = x_i$.*

Let us first show why this claim is true under the following idealized assumptions: (1) x is perfectly balanced and (2) the graphs $H^b_{a_1 a_2}$ are perfectly expanding in the sense that for every $i_1 \in S_{a_1}$ and subset of vertices $T \subseteq S_{a_2}$, i_1 has exactly $p^b_{a_1 a_2} D|T|$ edges going into T. Then the probability that in a random equation j, it happens that $x_{i_1} \oplus x'_{i_2} = f(x)_j$ is exactly $\Pr[z_1 \oplus z'_2 = P(z)]$, where $z = (z_1, \ldots, z_d) \in \{0,1\}^d$ is chosen uniformly at random and z'_2 is chosen from the distribution x'_{i_2} conditioned on $z_2 = x_{i_2}$ for a random $i_2 \in [n]$. It is easier to work with expectations instead of probabilities, so we consider the expression

$$\mathrm{E}[(-1)^{z_1 + z'_2 + P(z)}] = \mathrm{E}[(-1)^{z_1 + P(z)} \, \mathrm{E}[(-1)^{z'_2} \mid z_2]]$$

where $z = (z_1, \ldots, z_d)$. Taking the Fourier transform, we can write $\mathrm{E}[(-1)^{z'_2} \mid z_2]] = \alpha(-1)^{z_2} + \alpha'$, where $|\alpha'| \leq 1$. It follows that

$$\mathrm{E}[(-1)^{z_1 + z'_2 + P(z)}] = \mathrm{E}[\alpha(-1)^{z_1 + z_2 + P(z)} + \alpha'(-1)^{z_1 + P(z)}] \geq 2\alpha\gamma_2 - |\alpha'|\gamma_1 \geq \alpha\gamma_2$$

since by assumption $\gamma_1 \leq \alpha\gamma_2$. Therefore

$$\mathrm{E}[(-1)^{z_2'+P(z)} \mid z_1 = 0] - \mathrm{E}[(-1)^{z_2'+P(z)} \mid z_1 = 1] \geq 2\alpha\gamma_2$$

which we can rewrite as

$$\Pr[z_2' = P(0, z_2, \ldots, z_d)] - \Pr[z_2' = P(1, z_2, \ldots, z_d)] \geq \alpha\gamma_2 \tag{3.9}$$

and so the cases $x_{i_1} = 0$ and $x_{i_1} = 1$ can be distinguished by looking at the values $x'_{i_2} + P(x)_j$ over all edges $j = (i_1, i_2)$.

In the proof, we will replace each of these idealized assumptions with realistic counterparts that hold approximately and argue that the errors incurred by these approximations are not large.

Proof. We will show that because the graphs $H^b_{a_1 a_2}$ are expanding, for most $i_1 \in [n]$, the probability that $x_{i_1} = f(x)_j \oplus x'_{i_2}$ for a random equation $j = (i_1, i_2)$ that involves i_1 is close to the probability that $x_{i_1} = P(x_{i_1}, x_{i_2}, x_{i_3}, \ldots, x_{i_d}) \oplus x'_{i_2}$, where $i_2, i_3, \ldots, i_d$ are chosen uniformly at random from $[n]$. Then we will show that for all such i, $y_i^{(k)} = x_i$ for some $k \in [n]$.

Fix the assignment x and an index $i_1 \in [n]$ with $x_{i_1} = a_1$. We say an equation $u_{i_1} \oplus u_{i_2} = f(x)_j$ is of type (a_2, a'_2, b) if $x_{i_2} = a_2$, $x'_{i_2} = a'_{i_2}$, and $b = 1$ if $x_{i_1} \oplus x_{i_2} = f(x)_j$, and $b = 0$ if $x_{i_1} \oplus x_{i_2} \neq f(x)_j$. Let us say i_1 is *good* if for all types (a_2, a'_2, b), the number of equations of type (a_2, a'_2, b) is within δD (where $\delta = \alpha\gamma_2/12$) of the quantity

$$\begin{aligned} &\Pr[x_{i_2} = a_{i_2} \wedge x'_{i_2} = a'_{i_2} \wedge P(a_1, x_{i_2}, \ldots, x_{i_d}) = x_{i_1} \oplus x_{i_2}] \cdot D \quad \text{if } b = 1 \\ &\Pr[x_{i_2} = a_{i_2} \wedge x'_{i_2} = a'_{i_2} \wedge P(a_1, x_{i_2}, \ldots, x_{i_d}) \neq x_{i_1} \oplus x_{i_2}] \cdot D \quad \text{if } b = 0. \end{aligned}$$

In these probabilities, $i_2, \ldots, i_d$ are chosen uniformly at random from $[n]$. Let's assume that i_1 is good. Adding these probabilities over the relevant choices of a, a', b, we obtain that the number of equations j that contribute to Q_{i_1} is within 4δ of $p_{a_1} D$, where

$$p_{a_1} = \Pr_{i_2, \ldots, i_d}[x'_{i_2} = P(a_1, x_{i_2}, \ldots, x_{i_d})].$$

And so for every good vertex i_1, we have that $|Q_{i_1} - p_{a_1} D| \leq 4\delta$. Using the assumption that x is δ/d-balanced, we have that

$$|p_{a_1} - \Pr_{z_2, \ldots, z_d, z'_2}[z'_2 = P(a_1, z_2, \ldots, z_d)]| \leq \delta$$

where $z_2, \ldots, z_d$ are chosen uniformly at random from $\{0, 1\}$, and $z'_2 \in \{0, 1\}$ is a random variable chosen from the conditional distribution $\Pr[z'_2 = a_2 \mid z_2 = a_2] = \frac{1}{2}(1 + \alpha_{a_2})$. Using (3.9) and the triangle inequality, we can conclude that

$$p_0 - p_1 \geq \alpha\gamma_2 - 2\delta.$$

So there must be a difference of at least $(\alpha\gamma_2 - 10\delta)D > 0$ among those Q_{i_1} where $x_{i_1} = 0$ and those Q_{i_1} where $x_{i_1} = 1$, as long as i_1 is good. It follows that by choosing k appropriately, the assignment $y^{(k)}$ is correct on all good inputs. We now show that the number of good inputs is $n - O(\eta n/\delta)$.

To upper bound the number of inputs i_1 that are not good, we will bound this number for every choice of a_2, a_2', b and take a union bound. Since all the cases are analogous, for simplicity let us assume that $a_1 = 0$ and $a_2 = a_2' = 0, b = 1$. Let B_- (resp., B_+) be the set of i_1 with fewer (resp., more) than $\frac{1}{2}(1 + \alpha_0)p_{00}^1 D - \delta D/6$ neighbors i_2 in the graph H_{00}^1. By expansion, the number of edges (i_1, i_2) where $i_1 \in B_-$ and $x_{i_2} = x'_{i_2} = 0$ must be at least $\frac{1}{2}(1+\alpha_0)p_{00}^1 D|B_-| - \eta Dn$, so $\delta D|B_-|/6 \leq \eta Dn$, from where $|B_-| \leq 6\eta n/\delta$. By analogous reasoning, we can obtain the same upper bound on the size of the set B_+. Taking a union over all such sets for all choices of a_2, a_2', b, we conclude that the number of bad vertices is at most $O(\eta n/\delta)$. □

Proof of Proposition 3.4. We first argue that with high probability, a $1/2 + \gamma_2/4$ fraction of the equations (3.3) are satisfied by x. If x was perfectly balanced, then for every output of G the values of the inputs are chosen from the uniform distribution on $\{0,1\}^d$, and for every $j \in [m]$ we would have

$$\Pr[f_j(x) = x_{\Gamma(j,1)} \oplus x_{\Gamma(j,2)}] - \Pr[f_j(x) \neq x_{\Gamma(j,1)} \oplus x_{\Gamma(j,2)}] = \gamma_2.$$

When x is merely $\gamma_2/2d$-balanced, it still holds that

$$\Pr[f_j(x) = x_{\Gamma(j,1)} \oplus x_{\Gamma(j,2)}] - \Pr[f_j(x) \neq x_{\Gamma(j,1)} \oplus x_{\Gamma(j,2)}] \geq \gamma_2/2.$$

and so equation j is satisfied by x with probability $1/2 + \gamma_2/4$. Since the equations are independent, by the Chernoff bound, x satisfies at least a $1/2 + \gamma_2/8$ fraction of the equations with probability $1 - 2^{-\Omega(\gamma_2^2 m)} \geq 1 - 2^{-2n}$. Assume that x satisfies this many equations. By Theorem 3.2, x' then satisfies $1/2 + \Omega(\gamma_2/\log(1/\gamma_2))$ of the equations.

Now assume that the graphs H and $H_{aa'}^b$ are all $2/\sqrt{D}$-expanding, and each $H_{aa'}^b$ has $p_{aa'}^b m \pm 2m/\sqrt{D}$ edges. By Claim 3.5 and Claim 3.6, this happens with probability $1 - O(2^{-2n})$. By Claim 3.7, it follows that $|\alpha| \geq \sqrt{\gamma_2/\log(1/\gamma_2)}$. If α is positive, by Claim 3.8, we obtain that for some k, $y^{(k)}$ matches x $O(1/(\sqrt{D}\alpha\gamma_2))$ of the inputs i. If α is negative, we apply Claim 3.8 to the complementary assignment $\overline{x'}$ and obtain the conclusion for the assignment $\overline{y^{(k)}}$. □

4. From Almost Correct to Correct

In this section, we show that if we start with an almost correct assignment, $f_{G,P}(x)$ can be inverted for any nontrivial predicate P, provided that the constraint to variable ratio $m/n = D$ is a sufficiently large constant (depending on d).

PROPOSITION 4.1. *Let K be a sufficiently large constant, P be a non-constant predicate, and $r \geq 1$ be a parameter. Suppose $\eta \leq \beta^2/(Kd^6)$, $D \geq Kd^8/\beta^2$ and $r \leq \beta^K n/(Kd^K)$, and $d \leq \beta n^{1/K}/K$, where $\beta = \beta(P)$. There exists an algorithm that runs in time $D^3 n^{O(r)}$ such that for $1 - O(n^{-r})$ fraction of pairs (G, x), on input G, P, $f_{G,P}(x)$, and $x' \in \{0,1\}^n$ that has correlation $1 - \eta$ with x, the algorithm outputs an inverse for $f_{G,P}(x)$.*

The algorithm has three stages. In the first stage, the objective is to come up with an assignment that matches x on "core" inputs. Roughly speaking, the *core* of G with respect to the assignment x is the set of those inputs that are typical in the sense that they do not affect too many or too few constraints of G. The core of a random graph is likely to include most of the inputs. In the second stage, the algorithm unassigns some of the variables. At the end of this stage, there are no errors in the assignment, and all the inputs in the core are assigned (correctly). In the third stage, an assignment for the remaining variables is found by brute force. (The final assignment may not be x, as there are likely to be many possible inverses for $f_{G,P}(x)$.)

4.1. Support and Core. For $x \in \{0,1\}^n$ we write x^i for the string obtained by flipping the ith bit of x.

DEFINITION 4.2. *For $i \in [n], j \in [m]$, we say that the ith input* supports *the jth constraint with respect to an assignment $x \in \{0,1\}^n$ and graph G if $f_{G,P}(x^i)_j \neq f_{G,P}(x)_j$.*

We illustrate one role that the notion of support plays in the first stage of the algorithm. With high probability over the choice of planted assignment x and graph G, we expect most inputs of x to support a relatively large number of constraints. Suppose we are given an assignment y that is highly correlated but not identical to the planted assignment x. Since y is close to x, we expect most input bits of y to satisfy all the constraints they are involved in. So if an input i of y violates a noticeable number of constraints, we can view this as an indication that $x_i \neq y_i$ and flip its value with the hope of moving closer to x.

We would now like to argue that this procedure will affect all but a few exceptional inputs i where x_i and y_i differ. Suppose $y_i \neq x_i$, but y_i satisfies almost all the constraints it is involved in. If the ith input supports a noticeable number of constraints with respect to x, then there will also be a noticeable number of constraints that are satisfied by y but also supported by the ith input in x. Consider the input bits that participate in such a constraint. It cannot be the case that x_i and y_i differ only on the ith input bit, for otherwise we would get that $f_{G,P}(x^i)_j = f_{G,P}(x)_j$ (as both x and y satisfy the jth constraint). So there must be at least another input i' that participates in this constraint such that $x_{i'} \neq y_{i'}$. If this scenario happens too often (i.e., for too many values of i), using an expansion-like property of G we obtain a large set of inputs where x and y disagree, contradicting the assumption that x and y are highly correlated. By choosing parameters appropriately, it turns out that the number of disagreements between x and y drops by a factor of two, so after $\log n$ iterations all the disagreements vanish.

One implicit assumption we made in this argument is that the ith input participates in sufficiently many constraints. With high probability, this will be true for a random i, but we also expect to encounter some exceptions. Since these inputs are "atypical", we would like to discard them and deal with them separately later. Discarding a few atypical inputs (and the constraints they are involved in) may create more atypical ones. After discarding those and repeating sufficiently many times, we arrive at a set where every input is typical. We call this set the "core" of G.

Definition and properties of core Before we start the iteration process that arrives at the core, it will be convenient to discard some additional atypical inputs, such as those that support too few constraints. As we will use the construction several times with different parameters, we give a generic definition, which allows us to derive a core starting from any subset of the inputs.

Definition 4.3. *Let S be a subset of $[n]$. We say that a set H is an (S, λ, k)-*core of G *($\lambda \geq 0$, $k \geq 1$) if it can be obtained by the following iterative process:*

(i) $H_0 = A \cap S$, *where A is the set of inputs that appear in at least $(d-\lambda)D$ and at most $(d+\lambda)D$ constraints of G.*

(ii) If there exists an input $v_t \in H_t$ which appears in fewer than $(d - k\lambda)D$ constraints that contain only inputs from H_t, set $H_{t+1} = H_t \setminus \{v_t\}$,

(iii) If no such inputs exist at stage t, set $H = H_t$.

The construction of the core is nondeterministic: In step 2, there may be several available choices for inputs to be eliminated, and different choices of inputs may lead to different sets H.

For the definition of A in step 1, notice that on average, every input appears in dD constraints. Therefore, A captures those inputs whose appearance does not deviate much from their average. The following facts are easy consequences of the definition:

FACT 4.4. *Let H be a (S, λ, k)-core of G. If i is in H, then i appears in at most $(k+1)\lambda D$ constraints containing some input outside H.*

PROOF. Because i is in A, it appears in at most $(d+\lambda)D$ constraints. Because i survives the core elimination process, it appears in at least $(d-k\lambda)D$ constraints containing only inputs from H. So i can appear in at most $(k+1)\lambda D$ constraints containing some input outside the core. □

FACT 4.5. *If $S \subseteq S'$, then every (S, λ, k)-core of G is contained in every $(S', \lambda, k+2)$-core of G.*

PROOF. Let H and H' denote an (S, λ, k)-core of G and an $(S, \lambda, k+2)$-core of G, respectively. For contradiction, suppose that there exists some input i in H but not in H'. Consider the earliest stage t in the construction of H' where some $i \in H$ is eliminated from H'. Then $t > 0$ because initially H' contains all of H. But if i was eliminated at stage $t > 0$, then it appears in more than $(k+1)\lambda D$ constraints containing inputs i' that were eliminated at an earlier stage. Since all these inputs i' come from outside H, it follows that i appears in more than $(k+1)\lambda D$ constraints with some input outside H. This contradicts Fact 4.4. □

We now show that if a graph has a certain expansion-like property, then its core must be large. Using some standard probabilistic calculations, it will follow that a random graph is likely to have large core. We say G is $(\alpha, \alpha', \gamma)$-*sparse* if there do not exist sets V, V' of variables and C of constraints such that $|V| = \alpha n$, $|V'| = \alpha' n$, $|C| = \gamma D n$, and every constraint in C contains a pair of inputs $i \neq i'$ with $i \in V$ and $i' \in V'$. When $\alpha = \alpha'$, we say G is (α, γ)-sparse.

PROPOSITION 4.6. *Assume that $|A|, |S| \geq 1 - \varepsilon$ and G is $(3\varepsilon, \frac{2(k-1)\varepsilon\lambda}{d(d-1)})$-sparse. Then every (S, λ, k)-core of G has size at least $(1-3\varepsilon)n$.*

PROOF. We show that under the assumptions, the construction of the core can go through at most εn iterations. Since initially, $A \cap S$ has size $(1-2\varepsilon)n$ and one vertex is eliminated at every step, it follows that the core has size at least $(1-3\varepsilon)n$. We prove this by contradiction: If more than εn iterations are performed, G cannot be sparse.

Let $t > 0$. The input v_t (which was eliminated at stage t) appears in at least $(d-\lambda)D$ constraints. However, since v_t was eliminated at stage t, it can appear in at most $(d-k\lambda)D$ constraints containing only inputs from inside H_t. Let $T \geq t$ be any stage before the process terminates. It follows that $v_t \notin H_T$ must participate in at least $(k-1)\lambda D$ constraints that contain some (other) input from outside H_t, and therefore also outside H_T. Letting t range from 1 to T, it follows that there are at least $(k-1)\lambda DT$ pairs of inputs from outside H_T that appear in the same constraint. Each constraint can account for at most $\binom{d}{2}$ such pairs, so there are at least $(k-1)\lambda DT/\binom{d}{2}$ constraints that contain a pair of variables from outside H_T.

Now suppose for contradiction that the process takes more than εn steps. At stage $T = \varepsilon n$, we have $n - |H_T| \leq 3\varepsilon n$, but there are $(k-1)\lambda D\varepsilon n/\binom{d}{2}$ constraints that contain a pair of variables from outside H_T, contradicting of our assumption that G is $(3\varepsilon, \frac{2(k-1)\varepsilon\lambda}{d(d-1)})$-sparse. □

The core of a random graph We would now like to exclude those inputs that support too few constraints from the core. Let $\rho = \beta(P)/2^8$. In Proposition 4.8, we show that on average the ith input supports at least $32\rho D$ constraints with respect to x. Let A be the set of inputs that appear in no fewer than $(d-\rho)D$ and no more than $(d+\rho)D$ constraints. Let B be the set of inputs that support at least $30\rho D$ constraints with respect to x. We will use $H(G,x)$ to denote an arbitrary $(B,\rho,3)$-core of G.

By Proposition 4.6, to prove that $H(G,x)$ contains most of the inputs, it is sufficient to show that A and B are large and G is sparse. We begin by proving that a random G is likely to be sparse. We prove a slightly more general statement for later use.

PROPOSITION 4.7. *Assume $\alpha \leq \alpha'$ and $D \geq 4\alpha'/\gamma$. Then G is not (α,α',γ)-sparse with probability at most $(21d^4\alpha^2\alpha'/\gamma^2)^{\gamma Dn/2}$.*

PROOF. The probability that a specific constraint contains an input from V and an input from V' is at most $d^2\alpha\alpha'$ by a union bound. To upper bound the probability that G is not (α,α',γ)-sparse, we take a union bound over all triples

V, V', U of size αn, $\alpha' n$, and $\gamma D n$, respectively to obtain an upper bound of

$$\begin{aligned}
&\binom{n}{\alpha n}\binom{n}{\alpha' n}\binom{Dn}{\gamma Dn}\cdot(d^2\alpha\alpha')^{\gamma Dn} \\
&\qquad\le \Big(\frac{e}{\alpha'}\Big)^{2\alpha' n}\Big(\frac{e}{\gamma}\Big)^{\gamma Dn}(d^2\alpha\alpha')^{\gamma Dn} \\
&\qquad= \Big(\frac{e}{\alpha'}\Big)^{2\alpha' n}\Big(\frac{ed^2\alpha\alpha'}{\gamma}\Big)^{\gamma Dn} = \Big(\frac{e^2d^2\alpha}{\gamma}\Big)^{\gamma Dn}\Big(\frac{\alpha'}{e}\Big)^{(\gamma D-2\alpha')n} \\
&\qquad\le \Big(\frac{e^2d^2\alpha}{\gamma}\Big)^{\gamma Dn}\cdot\Big(\frac{\alpha'}{e}\Big)^{\gamma Dn/2} = \Big(\frac{e^3d^4\alpha^2\alpha'}{\gamma^2}\Big)^{\gamma Dn/2}
\end{aligned}$$

on the probability that G is not $(\alpha, \alpha', \gamma)$-sparse. □

We now prove that $H(G, x)$ is likely to be large. Our statement will be a bit more general as required for later application.

PROPOSITION 4.8. *For every pair of constants $a > 0, k > 1$ there exists a constant K such that the following holds. Assume that x is $\beta(P)/4d$-balanced. With probability $1 - 2^{-\Omega((\varepsilon\rho/d)\min\{\rho,1/d\}Dn)}$ over the choice of G, every $(B \cap J, a\rho, k)$-core of G has size at least $(1-3\varepsilon)n$, where $J \subseteq [n]$ is any set of size at least $(1-\varepsilon/2)n$, $\varepsilon \le \rho^2/Kd^8$, $n, D \ge Kd^2/\rho$.*

PROOF. By Proposition 4.6, the probability that $H(G, x)$ has size less than $1-3\varepsilon$ is at most the sum of the probabilities of the following three events: (1) $|A| < (1-\varepsilon)n$, (2) $|B| < (1-\varepsilon/2)n$, and (3) G is not $(3\varepsilon, \Omega(\rho\varepsilon/d^2))$-sparse. We now upper bound these probabilities.

We first upper bound the probabilities that $|A| < (1-\varepsilon)n$ and $|B| < (1-\varepsilon/2)n$. To bound $\Pr[|A| < (1-\varepsilon)n]$, we apply Lemma A.1 on the following bipartite graph: Vertices on the left come from $[n]$, vertices on the right come from $[Dn] \times [d]$, and $i \in [n]$ is connected to $(j, k) \in [Dn] \times [d]$ whenever $\Gamma(j, k) = i$ (namely, the ith input appears in position k in the jth constraint). By Lemma A.1, at most $\varepsilon n/2$ of the vertices on the left have fewer than $(1 - a\rho/d)dD$ or more than $(1 + a\rho/d)dD$ neighbors on the right with probability $1 - 2^{-\Omega((\rho^2\varepsilon/d)\cdot Dn)}$.

We now bound $\Pr[|B| < (1-\varepsilon/2)n]$. First, we lower bound the probability that the jth constraint is supported by at least one of its inputs. Let $z = (x_{\Gamma(j,1)}, \ldots, x_{\Gamma(j,d)})$. With probability $1 - d^2/2n \ge 1/2$, the inputs $\Gamma(j, k)$ are pairwise distinct for $1 \le k \le d$. Conditional on all inputs being pairwise distinct, each of the bits $z_1, \ldots, z_d$ is independent and at most $\beta/4d + d/n$-biased. By assumption, $\beta/4d + d/n \le \beta/2d$. Then the statistical distance

between the distribution on z and the uniform distribution is at most $\beta/2$. Under the uniform distribution, the boundary of P has probability β, so under the distribution on z it has probability at least $\beta/2$. Since the condition that $\Gamma(j,k)$ are pairwise distinct holds with probability $1/2$ or more, it follows that the jth constraint is supported by one of its inputs with probability at least $\beta/4 = 64\rho$.

Since the constraints are chosen independently, by a Chernoff bound with probability $1 - 2^{-\Omega(\rho Dn)}$, at least $32\rho Dn$ constraints are supported by at least one of their inputs. Consider one of these constraints. Conditional on the constraint being supported by one of its inputs, the first supporting input is distributed uniformly at random among all possible inputs. We can therefore apply Lemma A.1, which tells us that the probability of having more than $(\varepsilon/2)n$ constraints that support fewer than $30\rho D$ inputs is at most $1 - 2^{-\Omega(\rho^2\varepsilon Dn)}$.

By Proposition 4.7, the probability that G is not $(3\varepsilon, \Omega(\rho\varepsilon/d^2))$-sparse is at most $2^{-\Omega((\rho\varepsilon/d^2)Dn)}$. Adding all the failure probabilities, we obtain the desired bound. □

4.2. The Algorithm. To describe the algorithm, we need to introduce a bit more notation. Let $V \subseteq [n]$ be a collection of inputs and $C \subseteq [n]^d$ be a collection of constraints. Let $G_{V,C}$ be the bipartite graph with vertex set (V, C) and where an edge (i, j) is present whenever input i participates in constraint j. Recall that r is a parameter that controls the tradeoff between the running time and the success probability of the algorithm, and $\rho = \Omega(\beta(P))$, the boundary of P.

Algorithm Complete:
Input: A predicate P, a graph G, the value $f_{G,P}(x)$, an assignment $x' \in \{0,1\}^n$ (that correlates with x):

1. Set $\pi_0 = x$. For $k = 1$ to $\log n$ do the following: For each input i, if i appears in $5\rho D$ outputs unsatisfied by π_{k-1}, set $\pi_k(i) = \neg\pi_{k-1}(i)$. For others set $\pi_k(i) = \pi_{k-1}(i)$. Create all assignments y that differ from $\pi_{\log n}$ in at most r inputs.

2. For each assignment y produced in Stage 1, let B_y be the set of those inputs that support at least $26\rho D$ constraints with respect to y. Compute any $(B_y, \rho, 5)$-core H_y of G. For every subset $I \subseteq H_y$ of size r and every possible partial assignment $a \in \{0,1\}^I$, create the following assignment

$z \in \{0,1,\perp\}^n$:

$$z_i = \begin{cases} y_i, & \text{if } i \in H_y - I \\ a_i, & \text{if } i \in I \\ \perp, & \text{otherwise.} \end{cases}$$

3. For each assignment z produced in Stage 2, let Z be the subset of inputs i such that $z_i = \perp$, and W be the subset of constraints in G that contain at least one input from Z. If all connected components of $G_{Z,W}$ contain at most $r \log n$ inputs, exhaustively search for an assignment for the unassigned inputs that satisfy all the constraints in that component, and replace the unassigned components of z by this assignment.

4. If any of the assignments produced at this stage maps to $f_{G,P}(x)$ under f, output this assignment. Otherwise, fail.

We analyze the running time of Algorithm Complete. Stage 1 consists of $\log n$ iterations each taking time $O(Dn)$ after which a collection of at most $r\binom{n}{r}$ assignments are produced. For each output of Stage 1, Stage 2 consists of a core computation (which could take time $O(Dn^2)$), for which another set of $r\binom{n}{r}$ assignments is produced. In Stage 3 we perform an exhaustive search of assignments over at most Dn components of size $r \log n$ each, which can be done in time n^{r+1}. Stage 4 applies a computation of f on every candidate assignment that survives Step 3. It follows that the running time is $D^3 n^{O(r)}$.

4.3. The First Stage.

PROPOSITION 4.9. *Assume G is $(\alpha, \alpha', \rho\alpha'/d^2)$-sparse for all $\alpha \leq \alpha', \eta$ and $\alpha' \geq r/n$. Assume also that x and x' agree on at least a $(1-\eta)$-fraction of inputs. Then for at least one of the assignments y obtained in Stage 1, x and y agree on all inputs in every core $H(G,x)$.*

The proof relies on the following claim:

CLAIM 4.10. *Under the assumptions of Proposition 4.9, let B_k be the subset of $H(G,x)$ on which π_k and x disagree. Then $|B_k| < \max\{|B_{k-1}|/2, r\}$ for every $k > 0$.*

PROOF OF PROPOSITION 4.9. As k takes value at most $\log n$, by Claim 4.10, $|B_{\log n}| \leq r$. That is, $\pi_{\log n}$ and x take the same value on all but at most r inputs in $H(G,x)$. Thus trying all assignments of all possible subsets of size r, at least one y will match x everywhere on $H(G,x)$. □

PROOF OF CLAIM 4.10. Let $H = H(G, x)$. We will show that every input in B_k has at least ρD constraints that contain another input from B_{k-1}. We will then conclude that B_k cannot be too large because G is sparse.

Assume that $i \in B_k$ for some $k > 0$. We have two cases:

Case $i \in B_{k-1}$: Then $\pi_{k-1}(i) = \pi_k(i)$, so the assignment to input i was not flipped at stage k. Therefore i appears in at most $5\rho D$ constraints unsatisfied by π_{k-1}. Since i is in H, i supports at least $30\rho D$ constraints with respect to x. So i supports at least $25\rho D$ constraints (with respect to x) that are also satisfied by π_{k-1}. Since $\pi_k(i) \neq x_i$, each such constraint must contain some other input i' such that $\pi_{k-1}(i') \neq x_{i'}$. Furthermore, by Fact 4.4, i appears in at most $4\rho D$ constraints that contain some input not from H. So at least $21\rho D$ of the constraints that i appears in contain some other input from B_{k-1}.

Case $i \notin B_{k-1}$: Since $i \in B_k$, input i must participate in at least $5\rho D$ constraints unsatisfied by π_{k-1}. Since those constraints are satisfied by x, each of them must contain an input on which x and π_{k-1} disagree. Furthermore, by Fact 4.4, i appears in at most $4\rho D$ constraints with inputs not from H, so at least ρD of those constraints have some input from B_{k-1}.

In either situation, every input in B_k must appear in at least ρD constraints that contain some other input from B_{k-1}. This gives $\rho D|B_k|$ pairs of inputs from $B_{k-1} \times B_k$ that participate together in a constraint. So at least $\rho D|B_k|/\binom{d}{2}$ constraints contain a pair of inputs from $B_{k-1} \cup B_k$.

We now prove the claim. Assume for contradiction that $|B_k| \geq r$ and $|B_k| \geq |B_{k-1}|/2$ for some $k > 0$. Consider the smallest such k. Then $|B_{k-1}| \leq \eta n$. We consider two cases. If $|B_k| \geq |B_{k-1}|$ we contradict the assumption that G is sparse with $\alpha = |B_{k-1}|/n$ and $\alpha' = |B_k|/n$. If $|B_{k-1}|/2 \leq |B_k| < |B_{k-1}|$, we contradict the assumption that G is sparse with $\alpha = |B_{k-1}|/n$ and $\alpha' = 2|B_k|/n$. □

4.4. The Second Stage.

PROPOSITION 4.11. *Assume that G is $(3\alpha, 44\rho\alpha/d^2)$-sparse for all $r/n \leq \alpha \leq 3\varepsilon$, x and y agree on all inputs in $H(G, x)$, and $|H(G, x)| \leq 3\varepsilon n$. Then for at least one of the assignments z contained at the end of Stage 2, all inputs in $H(G, x)$ are assigned a $\{0, 1\}$ value in z and for every $i \in [n]$, either $z_i = x_i$ or $z_i = \perp$.*

PROOF. *All inputs in H are assigned:* By Fact 4.4, every $i \in H$ appears in at most $4\rho D$ constraints containing an input outside H. Therefore, i supports at least $26\rho D$ constraints with respect to x where all inputs are from H. Since

x and y match on all inputs appearing in these constraints, i supports at least $26\rho D$ constraints with respect to y, so $i \in B_y$. In particular, $A \cap B \subseteq A \cap B_y$. By Fact 4.5, $H \subseteq H_y$.

All assigned inputs are assigned correctly: Let F be the set of inputs $i \in H_y$ such that $x_i \neq y_i$. We will show that $|F| \leq r$. It follows that at least one of the assignments z has all its assigned inputs assigned as in x.

Let i be an input in F. By assumption, i is not in H. Since i is in H_y, it supports at least $26\rho D$ constraints with respect to y. By Fact 4.4, i supports at least $20\rho D$ constraints with respect to y containing only inputs from H_y. Consider any such constraint. This constraint must contain another input i' such that $x_{i'} \neq y_{i'}$. Then i' is not in H either, so i' is also in F. Summing up, we obtain at least $20\rho D|F|$ pairs of inputs in F that appear together in some constraint. So there are at least $20\rho D|F|/\binom{d}{2}$ constraints that contain pairs of variables from F. Since F does not intersect H, it has size at most $3\varepsilon n$. Since G is sparse, this is only possible if $|F| \leq r$. □

4.5. The Third Stage. The correctness of the third stage will follow from the next proposition. This proposition is analogous to Lemma 5 in Flaxman (2003) and Proposition 6 in Krivelevich & Vilenchik (2006), but our proof is somewhat simpler.

PROPOSITION 4.12. *Assume $D \geq d^7/\rho^2$ and $Kr(d/\rho)^K \leq n$ for a sufficiently large constant K. Let $\overline{H}$ denote the set of all inputs that do not appear in $H(G, x)$ and W denote the set of all constraints that contain at least one input from $\overline{H}$. Then with probability at most $4 \cdot 2^{-r}$ (over the choice of G and x), every connected component of $G_{\overline{H},W}$ has fewer than r inputs.*

To prove Proposition 4.12, we want to upper bound the probability that $G_{\overline{H},W}$ contains a connected component with r or more vertices. If this is the case, then this component must contain a subset that is "minimal" in the following sense:

DEFINITION 4.13. *Let $V \subseteq [n]$ be a collection of inputs and $C \subseteq [n]^d$ be a collection of distinct constraints. We say that C is a* minimal connected cover *of V if $G_{V,C}$ is connected, but $G_{V,C'}$ is not connected for every C' that is a strict subset of C.*

If $G_{\overline{H},W}$ contains a connected component with s or more inputs, then there exist subsets $V \subseteq \overline{H}$ and $C \subseteq W$ such that $r \leq |V| < r + d$ and C is a minimal connected cover of V. To obtain such a pair (V, C), we first remove enough

arbitrary constraints from W so that the number of inputs from $\overline{H}$ that remain in them is between r and $r+d$. This is always possible as every constraint in W contains between 1 and d inputs from $\overline{H}$. We let V be the set of remaining inputs from $\overline{H}$ that are present in the remaining constraints. We then possibly eliminate some additional constraints to obtain a minimal connected cover C of V.

Therefore the probability that $G_{\overline{H},W}$ contains a connected component with $\log n$ or more vertices is upper bounded by the probability that there exists a pair (V,C) such that: (1) $r \leq |V| < r+d$, (2) C is a minimal connected cover of V, (3) V is contained in $\overline{H}$, and (4) all constraints of C are present in G.

To prove Proposition 4.12, we first upper bound the probability (over the choice of G and x) that conditions (3) and (4) are satisfied for a particular pair (V,C). We then take a union bound over all pairs (V,C) that satisfy conditions (1) and (2). Let $|V| = v$ and $|C| = c$.

Facts about connected covers: We prove three useful facts about connected covers.

FACT 4.14. *Let C be a connected cover of V. The number of inputs that are not in V but participate in some constraint of C is at most $dc - (v+c) + 1$.*

PROOF. There are at most dc pairs (i,j) such that input $i \in [n]$ participates in constraint $j \in C$. Since $G_{V,C}$ is connected, it must contain at least $v+c-1$ edges. Each such edge gives a pair (i,j) with $i \in V$ and $j \in C$. So there can be at most $dc - (v+c) + 1$ pairs (i,j) with $i \notin V$ and $j \in C$. □

FACT 4.15. *Let C be a minimal connected cover of V. Then $|C| < |V|$.*

PROOF. Since C is a connected cover of V, the graph $G_{V,C}$ is connected. Let T be a spanning tree of $G_{V,C}$. The vertices of T come from $V \cup C$. Since T is a tree, it has more leaves than internal vertices. Suppose that $|C| \geq |V|$. Then T must contain at least one leaf c coming from C. Since c is a leaf, removing c from C does not disconnect T, and so $C - \{c\}$ is also a connected cover of V. Therefore C is not minimal. □

FACT 4.16. *Let C be a minimal connected cover of V. The number of inputs in V that participate in $2d$ or more constraints of C is at most $v/2$.*

PROOF. There are at most dc edges in $G_{V,C}$, By Fact 4.15, $v \geq c$, so the average degree of a vertex in V is at most d. By Markov's inequality, at most half the vertices have degree $2d$ or more. □

Bounding the probability for a specific pair (V, C): We fix a pair (V, C). Let R denote the (random) collection of all constraints that appear in G. Then for any x, the probability that G chosen from the distribution $\mathcal{G}_{n,m,d}$ satisfies conditions (3) and (4) is:

$$
\begin{aligned}
(4.17)\quad & \Pr_{G\sim\mathcal{G}_{n,m,d}}[C \subseteq R \text{ and } V \subseteq \overline{H}] \\
&= \Pr_{G\sim\mathcal{G}_{n,m,d}}[C \subseteq R \text{ and } V \cap H(G,x) = \varnothing] \\
&= \Pr_{G\sim\mathcal{G}_{n,m,d}}[C \subseteq R] \Pr_{G\sim\mathcal{G}_{n,m,d}}[V \cap H(G,x) = \varnothing \mid C \subseteq R] \\
&= \Pr_{G\sim\mathcal{G}_{n,m,d}}[C \subseteq R] \Pr_{G\sim\mathcal{G}_{n,m-c,d}}[V \cap H(G \cup C,x) = \varnothing] \\
&\leq \binom{m}{c}\left(\frac{c}{n^d}\right)^c \cdot \Pr_{G\sim\mathcal{G}_{n,m-c,d}}[V \cap H(G \cup C,x) = \varnothing].
\end{aligned}
$$

Here, $G \cup C$ denotes the constraint graph obtained by adjoining the constraints of C to those of G. To obtain the third equality, we observe that a uniformly random multiset of m constraints conditional on containing C can be obtained by choosing a multiset of $m - c$ constraints uniformly at random and taking the union with C. In the rest of this section, we will implicitly assume that G is chosen from the distribution $\mathcal{G}_{n,m-c,d}$. The last inequality follows by taking a union bound over all possible sets of c outputs where the constraints in C could occur.

Let $J \subseteq [n]$ be the set of inputs that appear in fewer than $2d$ constraints of C.

FACT 4.18. *Suppose $D \geq 2d/\lambda$. Every $(S \cap J, \lambda, 2)$-core of G is contained in every $(S, 2\lambda, 3)$ core of $G \cup C$.*

PROOF. This proof is analogous to the proof of Fact 4.5. Let H and H' denote an $(S \cap J, \lambda, 2)$-core of G and an $(S, 2\lambda, 3)$ core of $G \cup C$, respectively. For contradiction, suppose that there exists some input i in H but not in H'. Consider the earliest stage t in the construction of H' where some $i \in H$ is eliminated from H'.

We first argue that $t > 0$. For this it is sufficient to show that $A_{G,\lambda} \cap J \subseteq A_{G\cup C,2\lambda}$, where $A_{G,\lambda}$ is the set of inputs in G whose degrees are between $(d-\lambda)D$

and $(d+\lambda)D$. For any input $v \in A_{G,\lambda} \cap J$, its degree in $G \cup C$ is at least $(d-\lambda)D$ and at most $(d+\lambda)D + 2d \leq (d+2\lambda)D$, so it belongs to $A_{G\cup C,2\lambda}$.

Now suppose $i \in H$ was eliminated from H' at stage $t > 0$. Then it appears in at least $4\lambda D$ constraints of $G \cup C$ containing inputs i' that were eliminated at an earlier stage. Since all these inputs i' come from outside H, it follows that i appears in more than $4\lambda D$ constraints of $G \cup C$. Because i is in H and therefore in J, it can participate in at most $2d \leq \lambda D$ constraints of C. Therefore there are at least $3\lambda D$ constraints of G that contain i and another input from H. This contradicts Fact 4.4. □

We now define $H'(G, x)$ to be a *random* $(B \cap J, \rho/2, 2)$-core of G. By "random core" we mean that the selection of v_t in step 2 of the definition of core will be performed uniformly at random among all possible choices. We now upper bound the right side of inequality (4.17) for a random $x \sim \{0,1\}^n$ as follows:

$$\Pr_{G,x}[V \cap H(G \cup C, x) = \varnothing] \leq \Pr_{G,x,H'}[(J \cap V) \cap H'(G, x) = \varnothing] \qquad \text{(by Fact 4.18)}$$

$$\leq \Pr_{G,x,H'}[(J \cap V) \cap H'(G, x) = \varnothing \mid |H'| > (1-3\varepsilon)n] \tag{4.19}$$

$$+ \mathrm{E}_{H'} \Pr_{G,x}[|H'(G, x)| \leq (1-3\varepsilon)n]. \tag{4.20}$$

Let $\varepsilon = 1/(KdD) \leq \rho^2/(Kd^8)$. We now upper bound these two probabilities. By Proposition 4.8, probability (4.20) is at most $2^{-\Omega(\mathrm{poly}(\rho/d)Dn)}$. Probability (4.19) can be bounded using the following simple but important observation:

FACT 4.21. *Conditioned on $|H'| = h$, the set $H'(G, x)$ is uniformly distributed among all sets of size h in $[n] - J$.*

PROOF. Let Z, Z' be any two sets of size h in $[n] - J$. We show a probability-preserving bijection between the triples (G, x, H') such that $H'(G, x) = Z$ and those triples such that $H'(G, x) = Z'$. Let π be any permutation on $[n]$ that is invariant of J and maps Z to Z'. Then π induces a map between triples (G, x, H') by acting on the indices of x, the inputs of G, and the elements of B and the inputs v_t in the definition of H', respectively. Clearly π is probability preserving and if $H'(G, x) = Z$, then $\pi(H')(\pi(G), \pi(x)) = Z'$. It follows that Z' is at least as probable an outcome for $H'(G, x)$ as Z. By symmetry, they must have the same probability. □

Using Fact 4.21, we can bound expression (4.19) by

$$\Pr_{G,x}[(J \cap V) \cap H'(G,x) = \varnothing \mid |H'(G,x)| = h] \leq (1-h/n)^{|J \cap V|} \leq (1-h/n)^{|V|/2}$$

where the last inequality uses Fact 4.16, and so

$$\Pr_{G,x}[(J \cap V) \cap H(G,x) = \varnothing] \leq 2^{-\Omega(\mathrm{poly}(\rho/d)Dn)} + (3\varepsilon)^v \leq 2(3\varepsilon)^v,$$

because $\varepsilon = 1/(KdD)$, $v \leq r+d$, and $Kr(d/\rho)^K \leq n$.

The union bound: We now upper bound the probability that conditions (1)-(4) are satisfied by taking a union bound over all pairs (V, C) that satisfy conditions (1) and (2). To do so, we need an upper bound on the number of minimal connected covers C of V. We count as follows: First, each input in V can be assigned to one of c constraints in one of d positions in this constraint, giving $(cd)^v$ possible choices. By Fact 4.14 C contains at most $dc-(v+c)+1$ additional inputs coming from outside V. These can be assigned in $(n-|V|)^{dc-(v+c)+1} \leq n^{dc-(v+c)+1}$ possible ways. So the number of minimal connected covers of V is at most $(cd)^v n^{dc-(v+c)+1}$. We now take the desired union bound:

$$\begin{aligned}
&\sum_{v=r}^{r+d-1} \sum_{c=1}^{v-1} \underbrace{\binom{n}{v}}_{\text{choice of } V} \cdot \underbrace{(cd)^v \cdot n^{dc-(v+c)+1}}_{\text{choice of } C} \cdot \binom{m}{c} \left(\frac{c}{n^d}\right)^c \cdot 2(3\varepsilon)^v \\
&\qquad \leq \sum_{v,c} \left(\frac{en}{v}\right)^v \cdot (vd)^v \cdot n^{dc-(v+c)+1} \cdot \left(\frac{eDn}{c}\right)^c \left(\frac{c}{n^d}\right)^c \cdot 2(3\varepsilon)^v \\
&\qquad = \sum_{v,c} (ed)^v (eD)^c \cdot 2(3\varepsilon)^v \leq \sum_{v=r}^{r+d-1} (e^2 dD)^v \cdot 2(3\varepsilon)^v \leq 4 \cdot 2^{-r}.
\end{aligned}$$

The last inequality holds because $\varepsilon = 1/(KdD)$ for a sufficiently large constant K.

4.6. Proof of Proposition 4.1. To prove Proposition 4.1, we upper bound the failure probability of each of the three stages in Algorithm Complete. Let $H(G,x)$ be an arbitrary $(2\rho, 3)$ core of G. By Proposition 4.9, at the end of stage 1, x and some y agree on $H(G,x)$ unless G is not $(\alpha, \alpha', \rho\alpha'/d^2)$-sparse

for some $\alpha \le \eta$ and $\alpha' \ge \max\{\alpha, r/n\}$. By Proposition 4.7, this happens with probability at most

$$\sum_{a=1}^{\eta n} \sum_{a'=\max\{a,r\}}^{n} \Big(\frac{21d^6a^2}{\rho^2 a' n}\Big)^{\rho a' D/2d^2} \le \sum_{a=1}^{r} \sum_{a'=r}^{n} \Big(\frac{21d^6 a}{\rho^2 n}\Big)^{\rho a' D/2d^2} + \sum_{a=r+1}^{\eta n} \sum_{a'=a}^{n} \Big(\frac{21d^6 a}{\rho^2 n}\Big)^{\rho a' D/2d^2}$$
$$\le \sum_{a=1}^{r} 2\Big(\frac{21d^6 a}{\rho^2 n}\Big)^{\rho r D/2d^2} + 2\sum_{a=r+1}^{\eta n} \Big(\frac{21d^6 a}{\rho^2 n}\Big)^{\rho a D/2d^2}$$
$$\le 2r\Big(\frac{21d^6 r}{\rho^2 n}\Big)^{\rho r D/2d^2} + 2\eta n\Big(\frac{21d^6 \eta}{\rho^2}\Big)^{\rho \eta n D/2d^2} \le n^{-r}.$$

Let $\varepsilon = \rho^2/(Kd^8)$. Assuming x and y agree on $H(G,x)$, by Proposition 4.11, at the end of stage 2, all inputs in $H(G,x)$ are assigned a $\{0,1\}$ value in z and for every $i \in [n]$, either $z_i = x_i$ or $z_i = \perp$, unless $|H(G,x)| \le 3\varepsilon n$ or G is not $(\alpha, 44\rho\alpha/d^2)$-sparse for some $r/n \le \alpha \le 3\varepsilon$. By Proposition 4.8 the first event happens with probability at most $2^{-\mathrm{poly}(\rho/d)Dn}$. By Proposition 4.7, the second event happens with probability at most

$$\sum_{a=r}^{3\varepsilon n} \Big(\frac{d^8 a}{90\rho^2 n}\Big)^{22\rho a D/d^2} \le 3\varepsilon n \max\Big\{\Big(\frac{d^8 r}{90\rho^2 n}\Big)^{22\rho r D/d^2}, \Big(\frac{d^8 \varepsilon}{30\rho^2}\Big)^{66\rho\varepsilon D n/d^2}\Big\} \le n^{-r}.$$

Let Z be the set of inputs i such that $z_i = \perp$ and W be the set of constraints that contain at least one input from Z. Let W' be the set of constraints that contain at least one input from H. Assuming $Z \subseteq \overline{H}$, the connected components of $G_{Z,W}$ are contained in the connected components of $G_{\overline{H},W'}$. By Proposition 4.12, $G_{\overline{H},W'}$ has a component with $r \log n$ or more inputs with probability at most $4n^{-r}$. In such a case, Algorithm Complete outputs the desired assignment.

5. Amplifying Assignments

In this section we give the proof of Theorem 1.3. We may assume that the predicate P is not constant, for otherwise the function is trivially invertible. Theorem 1.3 is proved in two stages. First, in Proposition 5.1 we show that given an assignment x' that has correlation ε with x, it is possible to obtain an assignment w that agrees with x on most of the inputs. We then apply Proposition 4.1 with w as advice to complete the inversion.

Proposition 5.1. *Let K be a sufficiently large constant, P be any predicate and $D > 2^{Kd}/\varepsilon^{2d-2}$. There is an algorithm* Amplify *with running time polynomial in n, $1/\varepsilon$, and 2^d with the following property. With probability $1 - 2^{-Kd}$*

over the choice of G, for a $1 - 2^{-\Omega(\varepsilon^2 n)}$ fraction of assignments x and every assignment x' that has correlation ε with x, on input $f_{G,P}$, $f_{G,P}(x)$ and x', algorithm Amplify *outputs assignments $z_1, \ldots, z_{\mathrm{poly}(n)}$ so that at least one of them agrees with x on a $1 - 2^{-Kd}$ fraction of inputs.*

Theorem 1.3 follows by combining Proposition 5.1 and Proposition 4.1. We turn to proving Proposition 5.1, namely we describe and analyze Algorithm Amplify.

Algorithm Amplify takes advantage of the assignment x' to get empirical evidence about the values of each input value x_i in the hidden assignment. Without loss of generality, let us assume that $P(z)$ depends on its first variable z_1. Then the distributions $\mathcal{D}_0^{\mathrm{perfect}}$ and $\mathcal{D}_1^{\mathrm{perfect}}$ given by $(z_2, \ldots, z_d, P(0, z_2, \ldots, z_d))$ and $(z_2, \ldots, z_d, P(1, z_2, \ldots, z_d))$ (where $z_2, \ldots, z_d$ are uniformly random bits) will be statistically distinguishable with advantage at least $2^{-(d-1)}$. Let us now assume that x is perfectly balanced. Now consider an output $f(x)_j$ where i is the first variable with $x_i = b$ and consider the distribution $\mathcal{D}_b$ given by $(x'_{\Gamma(j,2)}, \ldots, x'_{\Gamma(j,d)}, P(b, x_{\Gamma(j,2)}, \ldots, x_{\Gamma(j,d)}))$. By the randomness of G, we can view $\mathcal{D}_b$ as a noisy variant of $\mathcal{D}_b^{\mathrm{perfect}}$, where the noise in each of the first $d-1$ components is independently chosen from the conditional distribution of $x'_{i'}$ given $x_{i'}$ for random i'. We will argue that if $\mathcal{D}_0^{\mathrm{perfect}}$ and $\mathcal{D}_1^{\mathrm{perfect}}$ are distinguishable, so are $\mathcal{D}_0$ and $\mathcal{D}_1$. By looking at all the neighbors j of input i and their values $(x'_{\Gamma(j,2)}, \ldots, x'_{\Gamma(j,d)}, f(x)_j)$, we collect empirical evidence whether they were drawn from $\mathcal{D}_0$ or from $\mathcal{D}_1$, allowing us to guess the value of x_i with high confidence.

Let us fix a pair of assignments x and x' with correlation ε. We consider the probability distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ described as follows:

$$\mathcal{D}_b : \text{Choose } i_2, \ldots, i_d \sim [n] \text{ and output } (x'_{i_2}, \ldots, x'_{i_d}, P(b, x_{i_2}, \ldots, x_{i_d})).$$

In Claim 5.2 we will prove that the distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ have noticeable statistical distance. We will also argue shortly that the two distributions are efficiently distinguishable given $\log n$ bits of advice (that depends on x). So by obtaining enough samples from the distribution $\mathcal{D}_{x_i}$, we can distinguish with high probability between the cases $x_i = 0$ and $x_i = 1$, and recover the value of the input x_i.

We now describe algorithm Amplify. The algorithm will need to compute the distributions $\mathcal{D}_0$ and $\mathcal{D}_1$. Since the algorithm does not have access to x, we describe these two distributions in an alternative way. Let $\mathcal{F}$ be the following distribution over $\{0,1\}^2$: First, choose $i \in [n]$ at random, then output the pair (x_i, x'_i). Then $\mathcal{F}$ can be described using $O(\log n)$ bits, since each value of $\mathcal{F}$

occurs with a probability that is a multiple of $1/n$. Let (a, a') denote a pair sampled from $\mathcal{F}$. The distribution $\mathcal{D}_b$ can then be described as follows:

1. Uniformly and independently sample pairs $(a_j, a'_j) \sim \mathcal{F}$ for $j = 2, \ldots, d$.
2. Output $(a'_2, \ldots, a'_d, P(b, a_2, \ldots, a_d))$.

Algorithm Amplify:
Inputs: A predicate P, a graph G, the value $y = f_{G,P}(x)$, $\varepsilon > 0$, an assignment $x' \in \{0,1\}^n$ that ε-correlates with x.

For every distribution $\mathcal{F}$ on $\{0,1\}^2$, where all probabilites are multiples of $1/n$, compute and output the following assignment $z_{\mathcal{F}}$:

1. Compute the distributions $\mathcal{D}_0$ and $\mathcal{D}_1$.
2. For every $i \in [n]$, compute the empirical distribution $\hat{\mathcal{D}}_i$ which consists of all samples of the form $(x'_{i_2}, \ldots, x'_{i_d}, y_j)$ for every constraint j of $f_{G,P}$ such that $\Gamma(j,1) = i$ and $\Gamma(j,k) = i_k$ for $2 \le k \le d$.
3. Set $z_{\mathcal{F},i} = b$ if $\hat{\mathcal{D}}_i$ is closer to $\mathcal{D}_b$ than to $\mathcal{D}_{1-b}$ in statistical distance.

It is easy to see that algorithm Amplify runs in time polynomial in n, $1/\varepsilon$, and 2^d. To argue its correctness, first we show (Claim 5.2) that the distributions $\mathcal{D}_0$ and $\mathcal{D}_1$ are at noticeable statistical distance. Then we show (Claim 5.6) that with high probability over G, for most i the distribution $\hat{\mathcal{D}}_i$ is statistically close to $\mathcal{D}_{x_i}$.

Claim 5.2. *Let K be a sufficiently large constant, P be any nonconstant predicate and x and x' be two assignments such that x is $\varepsilon/16$-balanced and x' has correlation ε with x. Then the statistical distance between $\mathcal{D}_0$ and $\mathcal{D}_1$ is at least $(\varepsilon^2/K)^{d-1}$.*

We observe that the distance can be as small as $\varepsilon^{-\Omega(d)}$, for example if P is the XOR predicate on d variables, x is any balanced assignment, and x' is an assignment that equals 1 on a $1 - \varepsilon$ fraction of inputs and 0 on the other inputs.

To give some intuition about the proof, consider the extreme case when $x' = x$. Because P is not constant, there must exist some setting for $a_2, \ldots, a_n$ such that $P(0, a_2, \ldots, a_n) \neq P(1, a_2, \ldots, a_n)$. Then the samples of the type $(a'_2, \ldots, a'_n, \star)$ are completely disjoint in $\mathcal{D}_0$ and $\mathcal{D}_1$, and the distributions can be distinguished on the samples of this type which occur with probability at least $2^{-\Omega(d)}$.

When $x' \neq x$, it is no more the case that for the proper choice of $a'_2, \ldots, a'_n$, the samples of the type $(a'_2, \ldots, a'_n, \star)$ are disjoint in the two distributions. However, we can still argue that the statistical distance between them is noticeable.

We will need a standard lemma about linear operators. For a simple proof, see for instance Theorem 4.3 in Stewart & Sun (1990).

LEMMA 5.3. *Let T be a linear operator from $\mathbf{R}^n$ to $\mathbf{R}^n$ and σ be the smallest of its singular values. Assume that $\sigma \neq 0$. Then for every $g \in \mathbf{R}^n$, $\|Tg\| \geq |\sigma| \cdot \|g\|$.*

PROOF OF CLAIM 5.2. We observe that

$$\Pr[a' = 0] \geq \varepsilon/2 \text{ and } \Pr[a' = 1] \geq \varepsilon/2, \quad \text{where } (a, a') \sim \mathcal{F}. \tag{5.4}$$

If this was not the case, for example $\Pr[a' = 0] < \varepsilon/2$, using the condition that x is $\varepsilon/16$ balanced we would have that $\Pr[a = 1] < 1/2 + \varepsilon/16$, and so $\Pr[a' = a] \leq \Pr[a' = 0] + \Pr[a = 1] < 1/2 + \varepsilon$, contradicting the fact that x' has correlation ε with x. Similarly we can rule out the possibility that $\Pr[a' = 1] < \varepsilon/2$.

It follows that the probability $\mathcal{F}^{d-1}(a'_2, \ldots, a'_d)$ of sampling $a'_2, \ldots, a'_d$ in $d-1$ independent copies of $\mathcal{F}$ must satisfy $\mathcal{F}^{d-1}(a'_2, \ldots, a'_d) \geq (\varepsilon/2)^{d-1}$ for every $a'_2, \ldots, a'_d$. The statistical distance $\mathrm{sd}(\mathcal{D}_0, \mathcal{D}_1)$ between $\mathcal{D}_0$ and $\mathcal{D}_1$ can now be lower bounded by:

$$\begin{aligned}
\mathrm{sd}(\mathcal{D}_0, \mathcal{D}_1) &= \sum_{\mathbf{a}' \in \{0,1\}^{d-1}} 2 \cdot \mathcal{F}^{d-1}(\mathbf{a}') \cdot \left| \mathrm{E}_{\mathcal{F}^{d-1}}[P(0, \mathbf{a}) - P(1, \mathbf{a}) \mid \mathbf{a}'] \right| \\
&\geq 2 \cdot (\varepsilon/2)^{d-1} \cdot \max_{\mathbf{a}'} \left| \mathrm{E}_{\mathcal{F}^{d-1}}[P(0, \mathbf{a}) - P(1, \mathbf{a}) \mid \mathbf{a}'] \right| \\
&\geq 2 \cdot (\varepsilon/4)^{d-1} \cdot \Big(\sum\nolimits_{\mathbf{a}' \in \{0,1\}} \mathrm{E}_{\mathcal{F}^{d-1}}[P(0, \mathbf{a}) - P(1, \mathbf{a}) \mid \mathbf{a}']^2 \Big)^{1/2}, \qquad (5.5)
\end{aligned}$$

where $\mathbf{a} = (a_2, \ldots, a_d)$, $\mathbf{a}' = (a'_2, \ldots, a'_d)$, and the conditional expectation $\mathrm{E}_{\mathcal{F}^{d-1}}[\ \cdot \mid \mathbf{a}']$ is taken over independent choices of $a_2, \ldots, a_d$ where each a_i is sampled from the distribution $\mathcal{F}$ conditioned on a'_i.

To lower bound (5.5) we define the linear operator T_{d-1} on the space of functions $g \colon \{0,1\}^{d-1} \to \mathbf{R}$ defined by

$$(T_{d-1} g)(a'_2, \ldots, a'_d) = \mathrm{E}_{\mathcal{F}^{d-1}}[g(a_2, \ldots, a_d) \mid a'_2, \ldots, a'_d].$$

With this notation, the expression (5.5) equals $2(\varepsilon/4)^{d-1} \|T_{d-1} g\|$, where g is the function $g(a_2, \ldots, a_d) = P(0, a_2, \ldots, a_d) - P(1, a_2, \ldots, a_d)$. We will shortly

prove that the smallest singular value of T_{d-1} is at least $(\varepsilon/32)^{d-1}$. Applying Lemma 5.3, we obtain that $\mathrm{sd}(\mathcal{D}_0, \mathcal{D}_1) \geq 2(\varepsilon^2/128)^{d-1}$.

We are left with showing that the smallest singular value of T_{d-1} is at least $(\varepsilon/32)^{d-1}$. The operator T_{d-1} is a $(d-1)$-wise tensor product of T_1: If $e_{b_2,\ldots,b_d} \colon \{0,1\}^{d-1} \to \{0,1\}$ is the point function such that $e_{b_2,\ldots,b_d}(a_2,\ldots,a_d) = 1$ if $a_i = b_i$ for all $2 \leq i \leq d$ and 0 otherwise, then we have the decomposition

$$(T_{d-1}e_{b_2,\ldots,b_d})(a'_2,\ldots,a'_d) = ((T_1 e_{b_2})(a'_2)) \cdots ((T_1 e_{b_d})(a'_d))$$

This follows from the independence of the samples $(a_2, a'_2), \ldots, (a_d, a'_d)$. Since the singular values of the tensor product of matrices are obtained by taking pairwise products of the singular values of the matrices in the tensor product, it follows that the smallest singular value of T_{d-1} is σ^{d-1}, where σ is the smallest singular value of T_1. We now lower bound this singular value.

Let M be a 2×2 matrix representation of the operator T_1. Then the entries of M are

$$M(c, c') = \Pr_{\mathcal{F}}[a = c \mid a' = c'] = p_{cc'}/(p_{0c'} + p_{1c'}).$$

where $p_{cc'}$ is the probability of the pair (c, c') in $\mathcal{F}$. The singular values σ, σ' of M, where $\sigma \leq \sigma'$, are the square roots of the eigenvalues of $M^T M$, so they satisfy the relations

$$\sigma^2 + \sigma'^2 = \mathrm{Tr}(M^T M)$$
$$\sigma^2 \sigma^2 = \det(M^T M) = \det(M)^2$$

from where $\sigma^2 \geq \det(M)^2 / \mathrm{Tr}(M^T M)$. Since M is a matrix of probabilities, $\mathrm{Tr}(M^T M) \leq 4$, so it remains to show that $|\det(M)| \geq \varepsilon/16$. Calculating $\det(M)$ we obtain

$$\det(M) = \frac{p_{00}p_{11} - p_{10}p_{01}}{(p_{00} + p_{10})(p_{01} + p_{11})} \geq p_{00}p_{11} - p_{10}p_{01}.$$

Without loss of generality let us assume $p_{10} + p_{11} \leq 1/2$. Then we can write

$$p_{00}p_{11} - p_{10}p_{01} = (p_{00} + p_{01})p_{11} - (p_{10} + p_{11})p_{01}.$$

Since x is $\varepsilon/16$-balanced, $1/2 - \varepsilon/16 \leq p_{00} + p_{01}, p_{10} + p_{11} \leq 1/2 + \varepsilon/16$, and we obtain that

$$|(p_{00}p_{11} - p_{10}p_{01}) - (p_{11} - p_{01})/2| \leq \varepsilon/16.$$

We now show that $|p_{11} - p_{01}| \geq \varepsilon/4$. Suppose this was not the case. Then

$$|p_{00} - p_{10}| \leq |(p_{00} + p_{01}) - (p_{10} + p_{11})| + |p_{11} - p_{01}| \leq 2 \cdot \varepsilon/16 + \varepsilon/4 < \varepsilon/2$$

and so $|p_{00} + p_{11} - p_{01} - p_{10}| < \varepsilon$, contradicting the assumption that x and x' are ε-correlated. It follows that $|\det(M)| \geq \varepsilon/16$, concluding the proof. $\square$

CLAIM 5.6. *Let K be a sufficiently large constant. Assume that $D > K2^{Kd}/\eta^2$. With probability at least $1 - 2^{-Kn}$ over the choice of G, for every pair of assignments x and x', for at least a $1 - 2^{-Kd}$ fraction of the inputs i, $\hat{\mathcal{D}}_i$ is at statistical distance at most η from $\mathcal{D}_{x_i}$, where $\hat{\mathcal{D}}_i$ is the distribution defined in step 2 of Algorithm Amplify.*

PROOF. Set $\delta = 2^{-Kd-1}$. Fix x and x'. Let S_i be the set of all outputs of $f_{G,P}$ whose first input is i. By Lemma A.1 (Appendix A), with probability $1 - 2^{-\Omega(\delta Dn)}$, all but δn of the sets S_i have size at least $D/2$. Fix i such that $|S_i| \geq D/2$. We now upper bound the probability that the statistical distance between $\hat{\mathcal{D}}_i$ and $\mathcal{D}_{x_i}$ is more than η. The distribution $\mathcal{D}_{x_i}$ has support size 2^d, so it is sufficient to upper bound the probability that probabilities of any outcome in the two distributions differs by more than $\eta/2^d$. By the Chernoff bound (applied to the sum of indicator variables that a given outcome ω is observed in each of the samples), this probability is at most $2^{-\Omega(\eta^2 D/4^d)}$. Taking a union bound over all 2^d outcomes and using the assumption $\eta^2 D > Kd4^d$, we conclude that the statistical distance between the two distributions is at most η with probability $1 - 2^{-\Omega(\eta^2 D/4^d)} > 1 - 2^{-(K+4)/\delta}$ (using the assumption $D > K2^{Kd}/\eta^2$). Since the events that the statistical distance between $\mathcal{D}_{x_i}$ and $\hat{\mathcal{D}}_i$ exceed η are independent over i (conditioned on the sets $S_1, \ldots, S_n$), by the union bound, the probability that this event happens for δn of those is such that $|S_i| \geq D/2$ is at most

$$\binom{n}{\delta n} \left(2^{-(K+4)/\varepsilon}\right)^{\varepsilon n} \leq 2^n \cdot 2^{-(K+4)n} = 2^{-(K+3)n}.$$

Therefore, with probability at least $1 - 2^{-\Omega(\delta Dn)} - 2^{-(K+3)n} \geq 1 - 2^{-(K+2)n}$, at least $(1 - 2^{-Kd})n$ of the pairs of distributions $(\mathcal{D}_{x_i}, \hat{\mathcal{D}}_i)$ are within statistical distance η. The claim follows by taking a union bound over all pairs of assignments (x, x'). □

PROOF OF PROPOSITION 5.1. Let K be a sufficiently large constant. By Claim 5.6 with $\eta = (\varepsilon^2/K)^{d-1}/2$, with probability at least $1 - 2^{-Kn}$ over the choice of G, for all pairs of inputs (x, x') and all but 2^{-Kd} fraction of the inputs i, the statistical distance between D_{x_i} and $\hat{\mathcal{D}}_i$ is at most η. Let G be such a graph, x' be any assignment, and x be any assignment that is $\varepsilon/16$-balanced. By the Chernoff bound, x is $\varepsilon/16$-balanced with probability $1 - 2^{-\Omega(\varepsilon^2 n)}$. By Claim 5.2, the statistical distance between $\mathcal{D}_0$ and $\mathcal{D}_1$ is at least 2η, so for all but a 2^{-Kd} fraction of inputs i, algorithm Amplify will set $z_{\mathcal{F},i} = x_i$. □

Acknowledgements

We thank Eyal Rozenman for useful discussions at the initial stages of this work and Benny Applebaum for suggesting the algorithm in Section 3.2 and sharing many other insights that led to a simplified and improved presentation. The authors' work was supported in part by the National Natural Science Foundation of China Grant 60553001, the National Basic Research Program of China Grants 2007CB807900, 2007CB807901, and Hong Kong RGC GRF grant 2150617. A preliminary version of this paper appeared in the Proceedings of the 13th International Workshop on Randomization and Computation (2009).

References

Noga Alon & Nabil Kahale (1997). A Spectral Technique for Coloring Random 3-Colorable Graphs. *SIAM J. Comp.* **26**(6), 1733–1748. ISSN 0097-5397.

Benny Applebaum, Boaz Barak & Avi Wigderson (2010). Public-key cryptography from different assumptions. In *STOC '10: Proceedings of the 42nd ACM symposium on Theory of computing*, 171–180. ACM, New York, NY, USA. ISBN 978-1-4503-0050-6.

Benny Applebaum, Yuval Ishai & Eyal Kushilevitz (2004). Cryptography in NC^0. In *Proceedings of the 45th Annual Symposium on Foundations of Computer Science*, 166–175.

Benny Applebaum, Yuval Ishai & Eyal Kushilevitz (2006). On Pseudorandom Generators with Linear Stretch in NC^0. In *Proceedings of the 10th International Workshop on Randomization and Computation (RANDOM 2006)*, 260–271.

Andrej Bogdanov & Youming Qiao (2009). On the Security of Goldreich's One-Way Function. In *Proceedings of the 13th International Workshop on Randomization and Computation (RANDOM)*, 392–405.

Moses Charikar & Anthony Wirth (2004). Maximizing Quadratic Programs: Extending Grothendieck's Inequality. In *Proceedings of the 45th Annual Symposium on Foundations of Computer Science*, 54–60.

James Cook, Omid Etesami, Rachel Miller & Luca Trevisan (2009). Goldreich's One-Way Function Candidate and Myopic Backtracking Algorithms. In *Proceedings of the 6th Theory of Cryptography Conference (TCC)*, 521–538.

Abraham Flaxman (2003). A spectral technique for random satisfiable 3CNF formulas. In *SODA '03: Proceedings of the fourteenth annual ACM-SIAM symposium*

on Discrete algorithms, 357–363. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA. ISBN 0-89871-538-5.

Michel X. Goemans & David P. Williamson (1995). Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems Using Semidefinite Programming. *J. ACM* **42**(6), 1115–1145.

Oded Goldreich (2000a). Candidate one-way functions based on expander graphs. Technical report, Electronic Colloquium on Computational Complexity (ECCC).

Oded Goldreich (2000b). *Foundations of Cryptography: Basic Tools*. Cambridge University Press, New York, NY, USA. ISBN 0-52-179172-3.

Michael Krivelevich & Dan Vilenchik (2006). Solving random satisfiable 3CNF formulas in expected polynomial time. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on discrete algorithms*, 454–463. ACM, New York, NY, USA. ISBN 0-89871-605-5.

Elchanan Mossel, Amir Shpilka & Luca Trevisan (2003). On ϵ-Biased Generators in NC^0. In *Proceedings of the 44th Annual Symposium on Foundations of Computer Science*, 136–145.

Jeanette P. Schmidt & Eli Shamir (1985). Component structure in the evolution of random hypergraphs. *Combinatorica* **5**(1), 81–94.

G. W. Stewart & Ji-guang Sun (1990). *Matrix Perturbation Theory*. Academic Press, Inc. ISBN 0-12-670230-6.

Danny Vilenchik (2007). It's all about the support: a new perspective on the satisfiability problem. *Journal on Satisfiability, Boolean Modeling, and Computation* **3**, 125–139.

A. A fact about sampling

We give a fact about sampling which we use throughout our analysis. In the random graph used in Goldreich's function with n inputs and $m = Dn$ outputs, for any fixed input position k of the predicate P, in expectation an input of $f_{G,P}$ appears in position k exactly D times. The following lemma shows that this is representative for most inputs. In the statement of the lemma, H represents the subgraph of G obtained by keeping only the kth incident edge of every output.

LEMMA A.1. *Fix $\varepsilon < 1/2, \eta < 1$ and suppose $D > (8/\eta^2)\log(1/\varepsilon)$. Let H be a random bipartite graph with n vertices on the left, Dn vertices on the right, and where each vertex on the right has exactly one neighbor on the left, chosen uniformly and independently at random. For a left vertex i, let N_i denote the number of its neighbors. Then with probability $1 - 2^{-\Omega(\eta^2 \varepsilon D n)}$, fewer than εn of the random variables N_i take value less than $(1-\eta)D$ (resp., more than $(1+\eta)D$).*

PROOF. Let I denote the set of those i such that $N_i < D/2$. By a union bound, the probability of $|I| \geq \varepsilon n$ is at most $\binom{n}{\varepsilon n}$ times the probability that $N_1, \ldots, N_{\varepsilon n} < (1-\eta)D$. Let $N = N_1 + \cdots + N_{\varepsilon n}$. Then $\Pr[N_1, \ldots, N_{\varepsilon n} < (1-\eta)D] \leq \Pr[N < (1-\eta)\varepsilon D n]$. Since N is a sum of Dn independent Bernoulli variables, each with probability ε, by a Chernoff bound we have $\Pr[N < (1-\eta)\varepsilon D n] \leq e^{-\eta^2 \varepsilon D n/3}$. Therefore the probability that fewer than εn of the N_i take value less than $(1-\eta)D$ is at most $\binom{n}{\varepsilon n} \cdot e^{-\eta^2 \varepsilon D n/3} = 2^{-\Omega(\eta^2 \varepsilon n)}$ (using the bound $\binom{n}{\varepsilon n} \leq 2^{2n\varepsilon \log(1/\varepsilon)}$ which holds for $\varepsilon < 1/2$ and sufficiently large n, together with the assumption $D > (8/\eta^2)\log(1/\varepsilon)$). The probability that more than εn of the N_i exceed $(1+\eta)D$ is calculated analogously. □

Manuscript received 1 October 2009

ANDREJ BOGDANOV
Department of CSE and ITCSC
Chinese University of Hong Kong
Shatin, N.T., Hong Kong

YOUMING QIAO
ITCS, Tsinghua University
Beijing 100084, China