
Boolean function analysis has become an indispensable tool in understanding the limits of approximation algorithms for NP-optimization problems. These are problems where good solutions may be hard to find, but once a solution is available its quality can be easily ascertained.

An important class of NP-optimization problems are *constraint satisfaction problems*. One famous example is maximum satisfiability of 3CNF clauses, or MAX-3SAT. In this problem we are given constraints of the form

$$\begin{aligned}x_1 \vee \bar{x}_2 \vee x_3 \\ \bar{x}_1 \vee x_2 \vee \bar{x}_3 \\ x_1 \vee x_2 \vee x_2 \\ \bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3.\end{aligned}$$

and want to find an assignment that simultaneously satisfies as many of them as possible. In this example setting x_1 to false (0) and x_2, x_3 to true (1) satisfies all four constraints.

Another example is maximum solvability of linear equations modulo 2 with three variables per equation, or MAX-3LIN. Instances of this problem look like this:

$$\begin{aligned}x_1 + x_2 + x_3 &= 1 \\ x_1 + x_2 + x_4 &= 0 \\ x_1 + x_3 + x_4 &= 1 \\ x_2 + x_3 + x_4 &= 1.\end{aligned}$$

Given such a system of equations, how many of them can we simultaneously satisfy? In this example you can see that we cannot satisfy all four – since the left hand sides add to zero, while the right hand sides add to one – but we can satisfy three out of the four, for example by setting $x_1 = x_2 = x_3 = x_4 = 1$.

Both MAX-3SAT and MAX-3LIN are special cases of constraint satisfaction problems over binary alphabet.

1 Constraint satisfaction problems

A *q-ary constraint satisfaction problem (qCSP)* over alphabet Σ is specified by a collection of variables x_1, \dots, x_n taking values in Σ and a collection of *constraints* $\phi_1, \dots, \phi_m: \Sigma^q \rightarrow \{0, 1\}$, where constraint ϕ_j is associated with a q -tuple of variables x_{j_1}, \dots, x_{j_q} . We say an assignment $\mathbf{x} = x_1 \dots x_n \in \Sigma^n$ satisfies constraint j if $\phi_j(x_{j_1}, \dots, x_{j_q}) = 1$. We say a q CSP instance is *satisfiable* if there exists an assignment that simultaneously satisfies all the constraints.

In MAX-3SAT $q = 3$, $\Sigma = \{0, 1\}$ and the constraints are disjunctions of literals. In MAX-3LIN, $q = 3$, $\Sigma = \{0, 1\}$, and the constraints are linear equations in three variables modulo 2.

In general we are often interested in the following kind of problem: Given a q CSP instance where the constraints are of a specific type, how should we go about finding an assignment that satisfies

as many of them as possible? We can always try brute-force search over all possible assignments, but this takes exponential time.

Is it possible to do better? In the case of MAX-3SAT, the theory of NP-completeness tells us that if we find an optimal solution substantially faster than by brute-force search, then we could do so for every problem in NP, which is unlikely to be possible. This is true even if the MAX-3SAT instance is completely consistent: Even if there exists an assignment that satisfies all the constraints, finding such an assignment would take an inordinate amount of time in the worst case.

What about MAX-3LIN? If all constraints are simultaneously satisfiable, then we can find a satisfying assignment, i.e. a solution to the system of equations, via linear algebra. But what if the equations are inconsistent? Can we still find an assignment that satisfies a large fraction of them? In expectation, a random assignment will satisfy half the equations. It turns out that this is essentially the best possible performance any efficient algorithm can guarantee in the worst case.

We will use the following terminology: We say a task is NP-hard if given an algorithm that achieves this task in time at most $t(n) \geq n$ on all instances of size n , for every NP problem there exists some other algorithm that solves all instances of size n in time $t(p(n))$ for some polynomial p (that depends on the problem).

Theorem 1 (Håstad). *For any constants $\eta, \varepsilon > 0$ the following task is NP-hard: Given a MAX-3LIN instance in which at least a $1 - \eta$ fraction of the constraints are simultaneously satisfiable, return an assignment that satisfies at least a $(1 + \varepsilon)/2$ fraction of them.*

Since every 3LIN constraint can be represented by four 3CNF constraints, it follows that the following task is also NP-hard: Given a MAX-3SAT instance in which at least a $1 - \eta$ fraction of the constraints are simultaneously satisfiable, return an assignment that satisfies at least a $(7 + \varepsilon)/8$ fraction of them.

The proof of Theorem 1 consists of several steps. The last step uses Fourier analysis. We will work out that part carefully but first let us give a rough sketch of what happens before Fourier analysis comes into play.

2 The PCP theorem and parallel repetition

The PCP theorem says that approximate optimization is hard in general, but does not give precise quantitative information about the parameters involved:

Theorem 2. *There exists an alphabet Σ and constants q and $\varepsilon > 0$ for which the following task is NP-hard: Given a satisfiable q CSP instance over Σ , find an assignment that satisfies at least a $1 - \varepsilon$ fraction of the constraints.*

Once we have this general form of the theorem, we can make some simplifying assumptions, for example that $q = 2$. Formally, we will reduce the q CSP instance Φ from the PCP theorem to a 2CSP instance Ψ which also satisfies the theorem and has some additional nice properties.

The instance Ψ has two kinds of variables: In addition to the variables x_1, \dots, x_n from Φ , it also has variables y_1, \dots, y_m , each taking value in Σ^q . When \mathbf{x} satisfies Φ , y_j is supposed to encode the restriction of \mathbf{x} on coordinates (j_1, \dots, j_q) .

The constraints of Ψ will encode two requirements: (1) that y_j satisfies ϕ_j and (2) that \mathbf{y} is consistent with \mathbf{x} (i.e. that the k -th coordinate of y_j is indeed equal to x_{j_k}). Formally, Ψ will have qm constraints $\psi_{jk}(y_j, x_{j_k})$, $1 \leq j \leq m, 1 \leq k \leq q$ where

$$\psi_{jk}(y_j, x_{j_k}) = \text{“}\phi_j(y_j) \text{ is true and the } k\text{th coordinate of } y_j \text{ equals } x_{j_k}\text{”}.$$

Now suppose we have an algorithm that, given a satisfiable instance Ψ , finds an assignment satisfying a $1 - \varepsilon/q$ fraction of constraints. We will use this algorithm to do the analogous thing for Φ . So suppose Φ is satisfiable. Then by construction, so is Ψ , so we can find an assignment (\mathbf{y}, \mathbf{x}) that violates at most ε/q fraction of the constraints ψ_{ji} . We claim that \mathbf{x} can then violate at most an ε fraction of the constraints ϕ_j . For if \mathbf{x} violates some constraint ϕ_j , then it must be that either x_{j_k} differs from the k th coordinate of y_j , in which case ψ_{ji} is violated, or if not then $\phi_j(y_j)$ must be false, so $\psi_{j1}, \dots, \psi_{jq}$ are all violated. So every constraint ϕ_j that is violated by \mathbf{x} yields at least one constraint ψ_{jk} that is violated by (\mathbf{y}, \mathbf{x}) .

This argument shows that without loss of generality, in Theorem 2 we may assume that $q = 2$ and the instance is of the same “type” as Ψ . Specifically, we may assume that:

1. The instance is *bipartite*: The variables come partitioned into two sets y_1, \dots, y_m and x_1, \dots, x_n so that the first variables in every constraint is some y_j and the second variable is some x_i ,
2. The constraints are *projections*: For every constraint $\psi_{ji}(y_j, x_i)$ and every assignment to y_j , there is at most one assignment $\pi_{ji}(y_j)$ to x_i that makes $\psi_{ji}(y_j, x_i)$ true. (For convenience we relabeled the constraint $\psi_{jk}(y_j, x_{j_k})$ to $\psi_{ji}(y_j, x_i)$.)

Parallel repetition In order to argue the hardness of MAX-3LIN, we need the following strengthening of the PCP theorem:

Theorem 3. *For every $\gamma > 0$ there exists an alphabet Σ such that the following task is NP-hard: Given a satisfiable 2CSP bipartite instance with projection constraints over Σ , find an assignment that satisfies at least a γ -fraction of the constraints.*

The main difference between this statement and the original PCP theorem is that the algorithm here is merely required to satisfy a small γ -fraction of the constraints, and not a $1 - \varepsilon$ fraction of them for some small ε . One transformation that allows us to go from the original version to this stronger version is *parallel repetition*.

Given a 2CSP instance Ψ , the t -fold parallel repetition Ψ^t of Ψ is the following 2CSP. If Ψ has variables $x_1, \dots, x_n, y_1, \dots, y_m$ taking values in Σ , then Ψ^t has n^t variables $x_{i_1 \dots i_t}$, $i_1, \dots, i_t \in [n]$ and m^t variables $y_{j_1 \dots j_t}$, where $j_1, \dots, j_t \in [m]$, taking values in Σ^t . The “intended assignment” to $x_{i_1 \dots i_t}$ is $(x_{i_1}, \dots, x_{i_t})$ and similarly for the y s.

The constraints of Ψ^t are as follows: For every t -tuple of constraints $\psi_{j_1 i_1}(y_{j_1}, x_{i_1}), \dots, \psi_{j_t i_t}(y_{j_t}, x_{i_t})$ of Ψ there is a constraint $\psi_{j_1 \dots j_t i_1 \dots i_t}(y_{j_1 \dots j_t}, x_{i_1 \dots i_t})$ which evaluates to true if $\psi_{j_k i_k}$ evaluates to true on the k th coordinates of $y_{j_1 \dots j_t}$ and $x_{i_1 \dots i_t}$ for all k between 1 and t .

By construction, if Ψ is satisfiable so is Ψ^t (because if (\mathbf{y}, \mathbf{x}) is a satisfying assignment for Ψ then the intended assignment induced by (\mathbf{y}, \mathbf{x}) is a satisfying assignment for Ψ^t). Now suppose an algorithm managed to find an assignment $(\mathbf{y}^t, \mathbf{x}^t)$ that satisfies a γ -fraction of the constraints of

Ψ^t . We would like to use this assignment to get an assignment (\mathbf{y}, \mathbf{x}) that satisfies a $1 - \varepsilon$ fraction of the constraints of Ψ . If $(\mathbf{y}^t, \mathbf{x}^t)$ was one of the intended assignments obtained from (\mathbf{y}, \mathbf{x}) , we could argue like this. Suppose (\mathbf{y}, \mathbf{x}) violated at least an ε fraction of the constraints $\psi_{ji}(y_j, x_i)$. Then a random constraint $\psi_{j_{i_1} \dots j_{i_t}}$ of Ψ^t is satisfied by $(\mathbf{y}^t, \mathbf{x}^t)$ if and only if all the constraints $\psi_{j_{i_1}}, \dots, \psi_{j_{i_t}}$ are satisfied simultaneously by (\mathbf{y}, \mathbf{x}) . Since these are independent this happens with probability at most $(1 - \varepsilon)^t < \gamma$ if t is a sufficiently large constant.

Unfortunately, there is no reason to assume that $(\mathbf{y}^t, \mathbf{x}^t)$ is one of the intended assignments and a much more elaborate analysis is required to complete the proof.

We will now show how to derive Theorem 1 from Theorem 3. The proof is by reduction: We start with the 2CSP instance Ψ from Theorem 3 and convert it into a MAX-3LIN instance Ξ with the following properties:

1. If Ψ is satisfiable, then a $1 - \eta$ fraction of the equations in Ξ are simultaneously satisfiable;
2. Every assignment that satisfies a $(1 + \varepsilon)/2$ fraction of the equations in Ξ can be converted into an assignment that satisfies at least a γ fraction of the equations in Ψ .

3 Encoding assignments

How do we construct a 3LIN instance Ξ from a generic 2CSP instance Ψ ? Intuitively, we want to encode the variables of Ψ using the variables in Ξ and represent the constraints of Ψ using 3LIN constraints. One immediate obstacle we face is that the variables of Ψ take value in some large alphabet Σ , while the variables of Ξ are $\{0, 1\}$ -valued.

Perhaps the simplest encoding we can try is to represent each variable x_i in Ψ by a collection of $\{0, 1\}$ -valued variables $X_i(1), \dots, X_i(|\Sigma|)$ in Ξ . The intended value of the variable $X_{i\sigma}$ is 1 if x_i takes value σ , and 0 otherwise. Similarly, we represent y_j by $Y_j(1), \dots, Y_j(|\Sigma|)$.

We now want to write constraints in Ξ to satisfy properties 1 and 2 above. Let us start by imposing the natural constraints on the variables of Ξ . There are two such types of constraints.

1. First, we should require that (Y, X) encodes a valid assignment, i.e. for every i , exactly one of the variables $X_i(1), \dots, X_i(|\Sigma|)$ should be set to 1, and similarly for the Y variables.
2. We should also require that the constraints of Ψ are satisfied: If $Y_j(\tau) = 1$ and $X_i(\sigma) = 1$ then the assignment $(y_j = \tau, x_i = \sigma)$ satisfies the constraint ψ_{ji} .

There are various ways of writing down such constraints using boolean formulas. However, after playing with different formulations for a while, you will see that one of two things usually goes wrong: Either the constraints you write involve too many variables, or if they involve up to 3 variables they are not robust: By violating a small number of constraints you can obtain assignments that satisfy almost all the constraints of Ξ , no matter whether Ψ was satisfiable or not. Fortunately, these issues pertain to the specific encoding we tried to use and not to the proof strategy itself.

Let us reformulate the proof strategy without making reference to the specific encoding. We take the 2CSP instance Ψ from Theorem 3 and reduce it to a 3LIN instance Ξ . The reduction will look as follows. We replace each variable y_j and x_i (that take values in some large alphabet Σ) by a

collection of variables $Y_j(1), \dots, Y_j(c)$ and $X_i(1), \dots, X_i(c)$ taking $\{0, 1\}$ -values. We will think of Y_j and X_i as strings in $\{0, 1\}^c$ that encode the values of y_j and x_i respectively. Then each constraint $\psi_{ji}(y_j, x_i)$ will be replaced by a collection of 3LIN constraints. These constraints are supposed to “check” the following properties of the assignment:

1. (Y, X) encodes a valid assignment: Specifically, X_i should look like a proper encoding of some value in Σ (the intended value for x_i), and similarly for Y_j .
2. The value encoded by Y_j and the value encoded by X_i satisfy the constraint ψ_{ji} .

To see how this may be possible, let us begin with the first requirement. Given some assignment to a variable x (in Σ), how can we come up with a boolean encoding X of x so that proper encodings are specified by 3LIN constraints? In fact we already did this: If we take X to be the Hadamard encoding of x , then the statement “ X is a codeword of the Hadamard code” can be represented by the linearity constraints $X(s) + X(t) + X(s + t) = 0$ for all s, t in the domain.

Now suppose we have a single projection constraint $\psi(y, x)$, where y and x take values in Σ . This means for every y , there is at most one value $x = \pi(y)$ which makes $\psi(y, x)$ true. We are given some encodings X of x and Y of y and we want to encode the statement $x = \pi(y)$ into a collection of 3LIN formulas. If X is the Hadamard encoding of x , then $X(s) = \ell_s(x)$, where ℓ_s is the linear function $\langle s, \cdot \rangle$. Similarly, if Y is the Hadamard encoding of y , then $Y(t) = \ell_t(y)$. How can we check that $x = \pi(y)$? Suppose we observe X at position s and Y at position t . We would expect to see the values $\ell_s(x)$ and $\ell_t(y)$. If we chose s and t so that ℓ_t and $\ell_s \circ \pi$ are the same function, then we could simply check that $X(s) = Y(t)$ and this would give us evidence that $x = \pi(y)$. Notice that the constraint $X(s) = Y(t)$, or $X(s) + Y(t) = 0$, is a 2LIN constraint.

Unfortunately, unless we are very lucky with π , $\ell_s \circ \pi$ will not be a linear function at all. This suggests that it may be helpful to extend the encoding X . The Hadamard encoding of x tells us the value of all linear functions at x , but it seems that we may also want to know the values of some nonlinear ones. As long as we write down the values of some nonlinear functions in the encoding, why not write down all of them?

4 The long code

The *(binary) long code* over message set Σ encodes a message $\sigma \in \Sigma$ by a string dict_σ in $\{1, -1\}^{2^{|\Sigma|}}$. Each position of the long code is indexed by a string $s \in \{0, 1\}^\Sigma$, and the encoding of a at position s is given by $\text{dict}_\sigma(s) = (-1)^{s\sigma}$. A corrupted codeword F could be any function $\{0, 1\}^\Sigma \rightarrow \{1, -1\}$, and the actual codewords are the *dictator functions* dict_σ .

A dictatorship test We now need a test for the long code based on 3LIN constraints. We have the linearity test as a starting point. This test always accepts all the dictator functions $F(s) = (-1)^{s\sigma}$, but unfortunately it also accepts all the other linear functions $\chi_a(s) = (-1)^{\langle a, s \rangle}$ for $|a| > 1$. Can we weed out those functions where $|a| > 1$?

We won’t quite achieve this, but here is one idea about how we can distinguish between those c that have small Hamming weight and those that have large Hamming weight. Let η be a small constant and choose a pair s, e from $\{0, 1\}^n$ independently by according to different distributions.

We choose s from the uniform distribution, while each coordinate of e is chosen from the η -biased distribution. This means each coordinate is chosen independently at random but takes value 1 with some small probability η and value 0 with probability $1 - \eta$.

Now consider the event $\chi_c(s) = \chi_c(s + e)$. The probability of this event is $\frac{1}{2}(1 + (1 - 2\eta)^{|c|})$. When $|c| = 1$ this probability is $1 - \eta$ which is close to one, but as $|c|$ becomes larger the probability approaches $1/2$ at an exponential rate. So if given χ_c we choose a random c and accept if $\chi_c(s) = \chi_c(s + e)$, we are much more likely to accept dictators than linear functions that depend on a lot of variables.

Now we combine this idea with the linearity test: Given a function $F: \{0, 1\}^n \rightarrow \{1, -1\}$, choose inputs s, t uniformly at random and e from the η -biased distribution and accept if $F(s)F(t)F(s + t + e) = 1$.

This test accepts all dictator functions with probability $1 - \eta$. Let's see what we can say about F if the test accepts it with probability at least $(1 + \varepsilon)/2$:

$$\begin{aligned} \varepsilon &\leq \mathbb{E}_{s,t,e}[F(s)F(t)F(s + t + e)] \\ &= \sum_{a,b,c} \hat{F}_a \hat{F}_b \hat{F}_c \mathbb{E}_{s,t,e}[\chi_a(s)\chi_b(t)\chi_c(s + t + e)] \\ &= \sum_{a,b,c} \hat{F}_a \hat{F}_b \hat{F}_c \mathbb{E}_s[\chi_{a+c}(s)] \mathbb{E}_t[\chi_{b+c}(t)] \mathbb{E}_e[\chi_c(e)] \\ &= \sum_a \hat{F}_a^3 \mathbb{E}_e[\chi_a(e)] = \sum_a \hat{F}_a^3 \prod_{i: a_i=1} \mathbb{E}_{e_i}[(-1)^{e_i}] = \sum_a \hat{F}_a^3 (1 - 2\eta)^{|a|}. \end{aligned}$$

Let k be the smallest integer so that $(1 - 2\eta)^{k+1} \leq \varepsilon/2$. Then by Parseval's identity

$$\sum_{a: |a|>k} \hat{F}_a^3 (1 - 2\eta)^{|a|} \leq \sum_{a: |a|>k} \hat{F}_a^2 (1 - 2\eta)^{|a|} \leq (1 - 2\eta)^{k+1} \leq \varepsilon/2$$

and so

$$\varepsilon/2 \leq \sum_{a: |a|\leq k} \hat{F}_a^3 (1 - 2\eta)^{|a|} \leq \sum_{a: |a|\leq k} \hat{F}_a^3 \leq \max_{a: |a|\leq k} \hat{F}_a$$

so F must have correlation $\varepsilon/2$ with some character χ_a with $|a| \leq k$. While we cannot say that such an F is close to a dictator, this weaker conclusion will be sufficient for what we need.

5 Hard instances of MAX-3LIN

Given a 2CSP instance Ψ as in Theorem 3, we now show how to construct a 3LIN instance Ξ as in Theorem 1. We start with a satisfiable instance Ψ , argue that in the corresponding Ξ at least a $1 - \eta$ fraction of constraints can be satisfied, apply an imaginary algorithm that finds an assignment satisfying at least $(1 + \varepsilon)/2$ of the constraints, and show how to convert it into an assignment that satisfies a γ -fraction of the constraints of Ψ , where $\gamma = \Omega(\eta\varepsilon^3/(\log 1/\varepsilon))$.

For each variable x_i of Ψ taking values in Σ , we introduce $2^{|\Sigma|-1}$ boolean variables X_i in Ξ . Similarly for every y_j we introduce such variables Y_j . The intended assignments to X_i and Y_j are the long code encodings of x_i and y_j , namely the truth-tables of the dictator functions dict_{x_i} and dict_{y_j} with one small modification.

For a reason that will become apparent later, it will be convenient to work with a slightly less redundant encoding. The dictator functions are odd: $\text{dict}_a(s) = -\text{dict}_a(\bar{s})$. So it is enough to specify the encodings $X_i(s)$ only for half of the inputs $s \in \{0, 1\}^\Sigma$; the value at the other inputs can be inferred using the formula $X_i(\bar{s}) = -X_i(s)$ and similarly for the Y_j s. This transformation is called *folding*.

We now describe the constraints of Ξ . We will specify what a random constraint of Ξ looks like. To get the instance consisting of all the constraints, we make a list of all possible random constraints (weighted by the probability of each constraint).

A random constraint of Ξ is created by the following experiment. We first choose a random constraint ψ_{ji} of Ψ . We now want a linear constraint that involves exactly three of the boolean variables among $X_i(s), Y_j(t)$ and “checks” that if Y_j and X_i were dictators, then the value encoded by Y_j projects to the value encoded by X_i according to the projection π_{ji} specified by ψ_{ji} .

To understand what this constraint should look like, suppose two of the variables involved in the constraint are $X_i(s)$ and $Y_j(t)$. What should the third one be? In the intended assignment, $X_i(s)$ and $Y_j(t)$ are the dictator functions s_{x_i} and t_{y_j} , where x_i and y_j satisfy the projection constraint $\pi_{ji}(y_j) = x_i$. Recall that we need to check three things: (1) that $X_i(s)$ looks like a dictator dict_σ ; (2) that $Y_j(t)$ looks like a dictator dict_τ and (3) that $\pi_{ji}(\tau) = \sigma$. To achieve this, we have the following three tests:

$$\begin{aligned} X_i(s)X_i(s')X_i(s + s' + e) &= 1 && \text{dictatorship test for } X_i \\ Y_j(t)Y_j(t')Y_j(t + t' + e) &= 1 && \text{dictatorship test for } Y_j \\ Y_j(s_{\pi_{ji}}) &= X_i(s) && \text{consistency test for } \pi_{ji}. \end{aligned}$$

In the last line, s_π is the string whose τ -th coordinate is $s_{\pi(\tau)}$, i.e. the $\pi(\tau)$ -th coordinate of s .

Doing these tests separately is wasteful. Instead we roll all three of them into one:

$$X_i(s)Y_j(t)Y_j(s_{\pi_{ji}} + t + e) = 1$$

where s and t are chosen uniformly at random, and n is chosen from the η -biased distribution.

Suppose Ψ has a satisfying assignment (\mathbf{x}, \mathbf{y}) . Let X_i and Y_j be the long code encodings dict_{x_i} and dict_{y_j} of the i th entry of \mathbf{x} and the j th entry of \mathbf{y} respectively. We now calculate the probability that a random constraint of Ξ is satisfied:

$$\begin{aligned} &\Pr_{i,j,s,t,e}[\text{dict}_{x_i}(s) \cdot \text{dict}_{y_j}(t) \cdot \text{dict}_{y_j}(s_{\pi_{ji}} + t + e) = 1] \\ &= \Pr_{i,j,s,t,e}[\text{dict}_{x_i}(s) \cdot \text{dict}_{y_j}(t) \cdot \text{dict}_{y_j}(s_{\pi_{ji}})\text{dict}_{y_j}(t)\text{dict}_{y_j}(e) = 1] \\ &= \Pr_{i,j,e}[\text{dict}_{y_j}(e) = 1] \\ &= 1 - \eta. \end{aligned}$$

If Theorem 1 was false, we would be able to efficiently find some other assignment $X_1, \dots, X_n, Y_1, \dots, Y_m$ that satisfies a $(1+\varepsilon)/2$ fraction of constraints of Ξ . We show how to use this assignment to produce a new one (\mathbf{x}, \mathbf{y}) that satisfies a γ -fraction of the constraints of Ψ .

Since $X_1, \dots, X_n, Y_1, \dots, Y_m$ satisfies a $(1 + \varepsilon)/2$ fraction of constraints of Ξ , we must have

$$\mathbb{E}_{i,j}[\mathbb{E}_{s,t,e}[X_i(s)Y_j(t)Y_j(s_{\pi_{ji}} + t + e)]] \geq \varepsilon$$

so by Markov's inequality, $\mathbb{E}_{s,t,e}[X_i(s)Y_j(t)Y_j(s_{\pi_{j_i}} + t + e)] \geq \varepsilon/2$ for at least $\varepsilon/2$ of the pairs (i, j) . Fix such a pair and to simplify notation let $X = X_i, Y = Y_j, \pi = \pi_{j_i}$. Applying Fourier expansion we get

$$\begin{aligned} \varepsilon/2 &\leq \mathbb{E}_{s,t,e}[X(s)Y(t)Y(s_{\pi} + t + e)] \\ &= \sum_{a,b,c} \hat{X}_a \hat{Y}_b \hat{Y}_c \mathbb{E}_{s,t,e}[\chi_a(s)\chi_b(t)\chi_c(s_{\pi} + t + e)] \\ &= \sum_{a,b,c} \hat{X}_a \hat{Y}_b \hat{Y}_c \mathbb{E}_s[\chi_a(s)\chi_c(s_{\pi})] \mathbb{E}_t[\chi_b(t)\chi_c(t)] \mathbb{E}_e[\chi_c(e)]. \end{aligned}$$

Clearly $\mathbb{E}_t[\chi_b(t)\chi_c(t)] = 1$ when $b = c$ and 0 otherwise and $\mathbb{E}_e[\chi_c(e)] = (1 - 2\eta)^{|c|}$. The term

$$\mathbb{E}_s[\chi_a(s)\chi_c(s_{\pi})] = \mathbb{E}_{e_s}[\langle a, s \rangle + \langle c, s_{\pi} \rangle] = \mathbb{E}_{e_s} \left[\sum_{\sigma} s_{\sigma} \left(a_{\sigma} + \sum_{\tau: \pi(\tau)=\sigma} c_{\tau} \right) \right].$$

(We use $\mathbb{E}e[\star]$ as shorthand for $\mathbb{E}[(-1)^{\star}]$.) Let $\text{odd}(c)_{\sigma} = \sum_{\tau: \pi(\tau)=\sigma} c_{\tau}$. This term vanishes unless $a = \text{odd}(c)$. So we get

$$\sum_c \hat{Y}_c^2 \hat{X}_{\text{odd}(c)} (1 - 2\eta)^{|c|} \geq \varepsilon/2.$$

Let k be the smallest integer so that $(1 - 2\eta)^{k+1} \leq \varepsilon/4$. Then by Parseval's identity

$$\sum_{c: |c| > k} \hat{Y}_c^2 \hat{X}_{\text{odd}(c)} (1 - 2\eta)^{|c|} \leq \varepsilon/4$$

and so

$$\begin{aligned} \varepsilon/4 &\leq \sum_{c: |c| \leq k} \hat{Y}_c^2 \hat{X}_{\text{odd}(c)} (1 - 2\eta)^{|c|} \\ &\leq \sum_{c: |c| \leq k} \hat{Y}_c^2 |\hat{X}_{\text{odd}(c)}| \\ &\leq \sqrt{\sum_{c: |c| \leq k} \hat{Y}_c^2} \sqrt{\sum_{c: |c| \leq k} \hat{Y}_c^2 \hat{X}_{\text{odd}(c)}^2} \\ &\leq \sqrt{\sum_{c: |c| \leq k} \hat{Y}_c^2 \hat{X}_{\text{odd}(c)}^2}. \end{aligned}$$

The second-to-last line follows by the Cauchy-Schwarz inequality and the last one uses Parseval's identity.

Now consider the following probabilistic algorithm for creating an assignment to Ψ : For every variable y_j , first choose $c \in \{0, 1\}^{\Sigma}$ with probability $\hat{Y}_{j,c}^2$, then choose τ uniformly at random from all τ such that $c_{\tau} = 1$ and assign τ to y_j . Similarly, for every x_i , choose $a \in \{0, 1\}^{\Sigma}$ with probability $\hat{X}_{i,a}^2$, then choose σ uniformly at random from all σ such that $a_{\sigma} = 1$ and assign σ to x_i .

(But what if we happened to choose $a = 0$ or $c = 0$ and no choice of σ and τ is possible? This cannot happen because of the folding. The folding guarantees that exactly half of the entries of X_i and Y_j are ones, and so $\hat{X}_{i,0}^2 = \hat{Y}_{j,0}^2 = 0$.)

We now argue that in expectation, this assignment satisfies at least a γ fraction of constraints of Ψ . Fix a “good” pair (i, j) for which $E_{s,t,e}[X(s)Y(t)Y(s \circ \pi + t + e)] \geq \varepsilon/2$. We will show that the constraint $\psi = \psi_{ji}$ is satisfied with probability at least $\varepsilon^2/16k$. By the calculation we just did

$$\sum_{c: |c| \leq k} \hat{Y}_c^2 \hat{X}_{\text{odd}(c)}^2 \geq \varepsilon^2/16$$

What is the probability that ψ is satisfied, i.e that $\pi(y) = x$ when x and y are chosen as above? Suppose we happened to choose c and $a = \text{odd}(c)$. Then there must exist a σ such that $a_\sigma = 1$. For every such σ , there must exist at least one τ such that $c_\tau = 1$ and $\pi(\tau) = \sigma$ (since the number of such y is odd). So conditioned on $a = \text{odd}(c)$, the probability of choosing a τ such that $\pi(\tau) = \sigma$ is at least $1/|c|$ and

$$\begin{aligned} \Pr[\pi(y) = x] &\geq \Pr[a = \text{odd}(c)] \cdot \Pr[\pi(y) = x \mid a = \text{odd}(c)] \\ &\geq \sum_c \hat{Y}_c^2 \hat{X}_{\text{odd}(c)}^2 \cdot \frac{1}{|c|} \geq \frac{1}{k} \sum_{c: |c| \leq k} \hat{Y}_c^2 \hat{X}_{\text{odd}(c)}^2 \geq \frac{\varepsilon^2}{16k} \end{aligned}$$

Since at least an $\varepsilon/2$ fraction of pairs (i, j) is good, in expectation the assignment will satisfy at least an $\varepsilon^3/16k$ fraction of the constraints.