

# Optimal Inapproximability with Universal Factor Graphs\*

Per Austrin, Jonah Brown-Cohen, and Johan Håstad

KTH Royal Institute of Technology

November 5, 2019

## Abstract

The factor graph of an instance of a constraint satisfaction problem (CSP) is the bipartite graph indicating which variables appear in each constraint. An instance of the CSP is given by the factor graph together with a list of which predicate is applied for each constraint. We establish that many Max-CSPs remains as hard to approximate as in the general case even when the factor graph is fixed (depending only on the size of the instance) and known in advance.

Examples of results obtained for this restricted setting are:

1. Optimal inapproximability for Max-3-Lin and Max-3-Sat (Håstad, J. ACM 2001).
2. Approximation resistance for predicates supporting pairwise independent subgroups (Chan, J. ACM 2016).
3. Hardness of the “ $(2 + \varepsilon)$ -Sat” problem and other Promise CSPs (Austrin et al., SIAM J. Comput. 2017).

The main technical tool used to establish these results is a new way of folding the long code which we call “functional folding”.

## 1 Introduction

Constraint Satisfaction Problems (CSPs), such as  $k$ -Lin and  $k$ -Sat, are some of the most well-studied problems in computational complexity. Already when considered as decision problems most CSPs are NP-complete and their maximization versions are hence NP-hard. In order to further investigate the computational complexity of these problems it is interesting to study restricted versions.

The factor graph of an instance of a CSP is the bipartite graph that connects  $u_i$  to  $v_j$  iff variable  $x_i$  appears in the  $j$ th constraint. The description of the instance is completed by indicating for each constraint which predicate is applied (in the case of e.g. Max-3-Sat this simply means specifying which variables are negated and which appear in their positive form). It is not difficult to find examples of factor graphs where the resulting Max-CSP instances are always easy; one example would be that the graph has bounded treewidth as such instances can be solved efficiently by dynamic programming. More generally, there has been extensive research on the tractability of CSPs with “left-hand side restrictions”, i.e., restrictions on the structure of the factor graph [DKV02, Gro07].

In this paper we are interested in the other end of the spectrum, namely to see if there is a sequence of factor graphs, one for each length of the input, such that the underlying Max-CSP remains hard if we restrict the instances to use these factor graphs. This general class of questions was first systematically considered by Feige and Jozeph [FJ12] who coined the term *universal factor graphs* for such sequences of factor graphs.

---

\*Research supported by the Approximability and Proof Complexity project funded by the Knut and Alice Wallenberg Foundation.

There are several reasons for studying this setting. On a fundamental level it is interesting to understand whether this separation into the factor graph and the predicate types applied can help us better understand the problem. In some situations one might have many instances with the same factor graph and it is interesting to study whether this is beneficial for a solver. This leads to the concept of “CSPs with preprocessing”. A particularly interesting case of this is for a linear problem such as Max- $k$ -Lin where the factor graph determines a fixed linear subspace, the negations give a point in space, and the problem is equivalent to finding a point in this fixed subspace that is close to the given point. This is a decoding problem for a linear code, but note that, in this case, the distances between the points supplied and the linear code is much larger than the minimal distance of the code. Hence, this problem is not really a traditional decoding problem for error correcting codes.

It is not surprising that some problems remain NP-hard in this setting. For instance, it is not difficult to see that if some basic parameters, such as the alphabet and the number of tapes for the involved Turing machine, are fixed then the standard proof of the Cook-Levin theorem produces an instance that can easily be made to have a universal factor graph.

Using a different approach Feige and Jozeph [FJ12] proved that universal factor graphs exist for the problem of deciding  $k$ -Sat. Furthermore, they also showed that the PCP Theorem can be established in the universal factor graph setting, opening up the possibility to proving hardness of approximation results for universal factor graphs. They went on to show that universal factor graphs exist for the problem of approximating Max- $k$ -Sat to within some constant factor, which for  $k = 3$  was  $77/80 + \varepsilon$ . Since general Max-3-Sat is well-known to be NP-hard to approximate within any constant larger than  $7/8 + \varepsilon$  [Hås01] this left open the problem whether fixing the factor graph might make Max-3-Sat easier to approximate.

In a follow-up work, Jozeph [Joz14] showed that there are universal factor graphs for every NP-hard Boolean CSP and for every APX-hard Boolean Max-CSP, but with possibly weaker approximation gaps than in the standard setting.

## 1.1 Our Results

We show that the hardness ratio of Max-3-Sat as well as those of many other problems carry over to the universal factor graph setting.

Essentially all our results are stated in the form of *approximation resistance*, i.e., that it is hard to approximate a certain Max-CSP significantly better than picking an assignment at random. In fact, in all cases but one, we prove *uselessness* in the sense of [AH13]. This means that it is NP-hard to find an assignment such that when picking a random constraint, the resulting distribution of  $k$ -tuples that are fed into the predicate deviates significantly from uniform, even under the promise that there exists an assignment that satisfies almost all or even all of the constraints.

We show that the following problems are all approximation resistant also in the universal factor graph setting.

1. The Max- $k$ -Sat problem on satisfiable instances for all  $k \geq 3$ .
2. The Max-TSA problem (each constraint is a Tri-Sum-And constraint, i.e., of the form  $x_1 + x_2 + x_3 + x_4 \cdot x_5 = b \pmod{2}$ ).
3. Any predicate supporting a pairwise independent subgroup such as linear equations modulo  $q$  and the very sparse “Hadamard predicate”.

Our proofs also give uselessness for all of these problems with the exception of Max-TSA. Approximation resistance for the Hadamard predicate is of particular interest because this gives optimal hardness of approximation for the general Max- $k$ -CSP problem (up to a constant factor independent of  $k$ ) [Cha16].

Hardness results for Max-TSA and similar problems are somewhat related to the one-way functions and pseudorandom generators proposed by Goldreich [Gol11]. For instance the Tri-Sum-And

predicate is one of the current standard candidate predicates for his construction, and our results show that there is a fixed set system such that it is NP-hard to distinguish strings that are in the image of the one-way function from strings that are at almost maximal distance from any output of the function. Of course in a cryptographic situation one is interested in an average case results while we are in a worst case setting. In any case we feel that our results give at least moral support to the conjecture that the functions proposed by Goldreich are indeed good pseudorandom generators.

We also consider hardness of Promise CSPs (PCSPs) as studied in several recent works [AGH17, BG18, BKO19]. In this setting we are given a satisfiable instance of some hard CSP (e.g. a 3-colorable graph) and are asked to find an assignment where each constraint is replaced by a weaker constraint (e.g. finding a  $c$ -coloring for some large  $c$ ). One example of this is the “ $(2+\varepsilon)$ -Sat” problem, in which we are given a  $(2k+1)$ -Sat instance where it is promised that there is an assignment that satisfies at least  $k$  literals in every clause, and the goal is to find an assignment that satisfies at least 1 literal in every clause. A fairly general hardness result of Brakensiek and Guruswami [BG18] states that if all the so-called polymorphisms of a PCSP satisfy a property called “ $C$ -fixing” then the PCSP is NP-hard. They further showed that this result is sufficient to establish a complete dichotomy for PCSPs where all constraints are symmetric functions. We were not able to obtain this result under universal factor graphs, but we can obtain an earlier result of Austrin et al. [AGH17], which applies if all the polymorphisms are  $C$ -juntas and the problem allows negations of variables. This in particular shows that the “ $(2+\varepsilon)$ -Sat” problem remains NP-hard in the universal factor graph setting.

Finally we also note that many typical gadget reductions carry over without modification to the universal factor graph setting and as a consequence our optimal hardness for Max-3-Lin also implies hardness of, e.g., 11/12-approximating Max-2-Lin and 21/22-approximating Max-2-Sat.

Our results generally apply to any domain, but to keep the notation as simple as possible in order to focus on the main new ideas, most parts of the paper present the details only for the Boolean case. We then briefly comment towards the end of the paper on the (minor) modifications needed for general domains.

## 1.2 Overview of Proof Techniques

To establish strong hardness results in the universal factor graph framework the key is to combine the universal factor graph property with a, at least somewhat, more modern soundness analysis of the underlying PCPs. The results of Feige and Jozeph [FJ12] essentially used the technique of Bellare et al. [BGS98] and was for instance unable to adapt the reductions of [Hås01]. The main technical problem is that the used folding of the long code (called conditioning) is difficult to combine with having a universal factor graph.

We circumvent this problem by introducing a new folding of the long code which we call “functional folding”. This folding is not as powerful as conditioning but sufficient for our purposes. Functional folding is based on the concept of “equational CSPs”. These are CSPs in which each constraint is of the form  $f(\mathbf{x}_S) = b$  where  $f$  is a function of constant arity and  $\mathbf{x}_S$  is the values of the variables from a small set  $S$ . Here, we demand that the only data that varies with the instance are the right hand sides  $b$ , while  $f$  and  $S$  remain fixed. Any such CSP which is APX-hard with a universal factor graph on satisfiable instances can be used as starting point (for concreteness we use the Max-TSA problem). Let us, without actually describing the details of what functional folding is (which the reader who wishes to skip ahead can find in Section 3), briefly explain what differences it results in compared to previous results.

We follow the standard “parallel repetition + long coding” reduction paradigm underlying the majority of strong inapproximability results. In this paradigm, the main objects in the reductions are Boolean functions  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  where each coordinate  $i \in [n]$  represents an assignment  $\mathbf{x}_T$  to some subset  $T$  of variables of the initial equational CSP. The “conditioning” type of folding normally employed in this setting further ensures that all of the  $n$  coordinates represent assignments  $\mathbf{x}_T$  that actually satisfy the equations on  $T$  in the initial CSP. The key difference in our setting where we apply functional folding instead of conditioning is that some of the coordinates now represent

assignments that do not satisfy the equations on  $T$ . Exactly which of the coordinates are satisfying depends (only) on the right-hand sides  $b$  of the underlying equations. Nevertheless, the functional folding still gives us a (weaker) guarantee (formalized in [Lemma 3.4](#)), which implies that for any non-zero Fourier coefficient of  $f$ , at least one of the associated variables represents a satisfying assignment.

Having only this weaker guarantee, it turns out to be possible to reprove many earlier results in the universal factor graph setting after some modifications in the constructions and the proofs. The most important modification is that, even in the easiest cases of proving that Max-3-Lin is NP-hard to approximate within  $\frac{1}{2} + \varepsilon$  or that  $(2 + \varepsilon)$ -Sat is NP-hard, we need to use a *smooth* parallel repetition as the starting point. This allows us to cope with the issue that only some of the coordinates of each purported long code are satisfying assignments (i.e. feasible assignments in the parallel repeated game).

Another difference, albeit a purely notational one, is the following. Because of the property of functional folding that only some of the coordinates correspond to feasible assignments in the parallel repeated game, we need to keep track for each long code of how the coordinates represent partial assignments to the underlying equational CSP. Due to this, we cannot view the parallel repeated game as an abstract label cover instance and in the process we lose most of the simplified notation afforded by this “modern” view. For example we have to view the purported long codes given to the PCP verifier as Boolean functions of Boolean functions rather than simply Boolean functions of generic bit strings.

The fact that we need a smooth parallel repetition implies that if we look more closely at the parameters of our results we do not get as strong results as in the general case. In the black and white world of polynomial vs non-polynomial time we match the general results, but in more fine-grained measures our results are weaker than the general case. Here we are referring to (i) the question of how small  $\varepsilon$  can be as a function of  $n$ , and (ii) the question of what quantitative lower bounds can be given on the running time assuming the Exponential Time Hypothesis. For these questions, our results for universal factor graphs do not match the state of art for the general case. We elaborate a bit more on this in the concluding remarks towards the end of the paper.

### 1.3 Organization

An outline of the paper is as follows. In [Section 2](#) we cover some background material. In [Section 3](#) we define our new notion of functional folding and set up the general framework for our hardness reductions. [Section 4](#) is devoted to the reproving the results of Håstad [[Hås01](#)] for basic problems such as Max-3-Lin and Max-3-Sat (and also Max-TSA for which the hardness is proved in the same way). In [Section 5](#) we show how to adapt the results of Chan [[Cha16](#)], and in [Section 6](#) we give results for Promise CSPs. Then in [Section 7](#) we discuss small extensions such as gadget reductions and larger domains, and finally in [Section 8](#) we round off with some remaining open questions.

## 2 Preliminaries

For a vector  $\mathbf{v} \in \Sigma^I$  indexed by  $I$  over some set  $\Sigma$ , and a vector  $S = (i_1, \dots, i_k) \in I^k$  of indices, we write  $\mathbf{v}_S$  for the vector  $(v_{i_1}, \dots, v_{i_k})$ . For a set  $\beta$  of vectors we let  $\beta_S = \{\mathbf{v}_S \mid \mathbf{v} \in \beta\}$  (viewed as a set and not a multiset).

For any event  $E(x)$  we let  $\mathbf{1}(E(x))$  denote the function that is one exactly when  $E(x)$  is true and zero otherwise. In a similar vein for a set  $\beta \subseteq I$  of indices  $\mathbf{1}_\beta \in \{0, 1\}^I$  is the vector that is one exactly on the set  $\beta$  and zero outside this set. If  $\beta$  is a single index, this is a unit vector.

For an index set  $I$  of coordinates we write  $\mathcal{F}_I = \{f : \{0, 1\}^I \rightarrow \{0, 1\}\}$  for the set of all Boolean functions on  $I$ , and for an integer  $n$  we write  $\mathcal{F}_n = \mathcal{F}_{[n]}$ .

When arguing about Boolean functions we let “+” denote exclusive-or. We also have addition of real numbers but hopefully the meaning of each + is clear from the context.

## 2.1 Constraint Satisfaction Problems

We begin by introducing notation for CSPs. As this paper focuses on the case of Boolean inputs we only give formal definitions in this special case. We expect that the reader is able to guess the more general definitions needed for our brief discussion of larger domains in [Section 7.2](#).

**Definition 2.1.** A  $k$ -ary constraint language  $\Gamma$  is a finite set of functions  $f : \{0, 1\}^k \rightarrow \{0, 1\}$ .

Different constraint languages give rise to different CSPs as formalized in the following definition.

**Definition 2.2.** An instance of the Max-CSP( $\Gamma$ ) problem is of the form  $I = (X, C)$  where  $X$  is a set of  $n$  variables and  $C$  is a set of  $m$  constraints. Each constraint  $c \in C$  is a pair  $c = (f, S)$  where  $S \in X^k$  (the *scope* of  $c$ ) is a tuple of  $k$  distinct variables and  $f \in \Gamma$  (the *constraint type* of  $c$ ). An assignment  $\mathbf{a} \in \{0, 1\}^X$  of values to the variables satisfies the constraint  $c = (f, S)$  if  $f(\mathbf{a}_S) = 1$ .

The objective is to find an assignment to the variables satisfying as many constraints as possible. An instance of Max-CSP( $\Gamma$ ) is  $\alpha$ -satisfiable if there is an assignment that satisfies an  $\alpha$  fraction of the constraints.

For every Max-CSP and some approximation parameters, there is a naturally associated approximation problem, which we generally phrase as the following promise decision problem.

**Definition 2.3.** For parameters  $0 \leq s \leq c \leq 1$ , the problem of  $(c, s)$ -approximating Max-CSP( $\Gamma$ ) is the promise problem where the goal is to distinguish between instances of Max-CSP( $\Gamma$ ) that are at least  $c$ -satisfiable, and those that are at most  $s$ -satisfiable.

For example, Max-3-Lin is the problem Max-CSP( $\Gamma$ ) where  $\Gamma$  consists of the two functions  $f_0(x, y, z) = x + y + z$  and  $f_1(x, y, z) = x + y + z + 1$ . In fact Max-3-Lin is a prototypical example of the following important subclass of CSPs which plays a critical role in our reductions.

**Definition 2.4.** An *equational CSP* is a CSP where the constraint language  $\Gamma$  contains exactly two functions:  $f$  and  $\neg f$  for some  $f : \{0, 1\}^k \rightarrow \{0, 1\}$ . Equivalently, an equational CSP is one where each constraint  $c$  has the form either  $f(\mathbf{x}_S) = 0$  or  $f(\mathbf{x}_S) = 1$ .

**Definition 2.5.** For a predicate  $f : \{0, 1\}^k \rightarrow \{0, 1\}$ , Max-CSP( $f^\pm$ ) is the CSP where each constraint is of the form  $f(\mathbf{x}_S + \mathbf{b}) = 1$  for some scope  $S$  and vector  $\mathbf{b} \in \{0, 1\}^k$ .

We now formally define the CSPs of interest in this paper.

**Definition 2.6.** The Tri-Sum-And predicate  $f_{\text{TSA}} : \{0, 1\}^5 \rightarrow \{0, 1\}$  is given by

$$f_{\text{TSA}}(x, y, z, u, v) = x + y + z + uv.$$

We write Max-TSA for the (equational) Max-CSP( $\{f_{\text{TSA}}, \neg f_{\text{TSA}}\}$ ) problem which is a well known NP-hard problem.

**Definition 2.7.** Max- $k$ -Sat is the Max-CSP( $\vee_k^\pm$ ) problem, where  $\vee_k : \{0, 1\}^k \rightarrow \{0, 1\}$  is the OR function on  $k$  Boolean variables.

**Definition 2.8.** Max- $k$ -Lin is the (equational) Max-CSP( $\{+_k, \neg +_k\}$ ) problem, where  $+_k : \{0, 1\}^k \rightarrow \{0, 1\}$  is the parity function on  $k$  Boolean variables.

**Definition 2.9.** For  $k$  of the form  $2^\ell - 1$  for some integer  $\ell$ ,  $\text{Had}_k : \{0, 1\}^k \rightarrow \{0, 1\}$  is the predicate where the indices are in one-to-one correspondence with the nonempty subsets of  $[\ell]$  and a string  $\mathbf{x}$  satisfies  $\text{Had}_k$  iff  $x_\alpha + x_\beta = x_\gamma$  whenever  $\alpha \Delta \beta = \gamma$ .<sup>1</sup>

It is easy to see that an accepted  $\mathbf{x}$  is determined by the  $\ell$  singleton variables  $x_{\{i\}}$  and that  $\text{Had}_k$  accepts exactly  $2^\ell = k + 1$  strings. We write Max-Had $_k$  for the Max-CSP( $\text{Had}_k^\pm$ ) problem.

---

<sup>1</sup>Here  $\alpha \Delta \beta$  is the symmetric difference of the two sets  $\alpha$  and  $\beta$ , i.e., the set of elements that appear in exactly one of the two sets.

## 2.2 Factor Graphs and Preprocessing

Next we define the factor graph of a CSP.

**Definition 2.10.** The factor graph of an instance  $I = (X, C)$  of  $\text{Max-CSP}(\Gamma)$  is the bipartite graph  $G = (X, Y, E)$ , where  $Y = \{S \mid (f, S) \in C\}$  is the multiset of scopes of the constraints of  $I$ , and there is an edge between a variable  $x \in X$  and scope  $S \in Y$  whenever  $x \in S$ .

Note that the factor graph precisely describes the scopes of the constraints of  $I$  but is independent of the constraint types  $f$  used for each constraint. Our results are about hardness of approximation for Max-CSPs where the factor graph is fixed in advance and the instance only consists of the constraint types. To this end, we make the following definitions.

**Definition 2.11.** A family  $\mathcal{F} = \{\mathcal{F}_n\}_{n>0}$  of factor graphs parameterized by  $n$  is *explicit* if  $\mathcal{F}_n$  can be constructed in time  $\text{poly}(n)$ .

**Definition 2.12.** We say that  $\text{Max-CSP}(\Gamma)$  is  $(c, s)$ -UFG-NP-hard if there is an explicit family  $\{\mathcal{F}_n\}$  of factor graphs and a polynomial time reduction  $R$  from 3-Sat instances  $I$  on  $n$  variables to  $\text{Max-CSP}(\Gamma)$  instances  $R(I)$  having factor graph  $\mathcal{F}_n$  such that

1. If  $I$  is satisfiable then  $R(I)$  is  $c$ -satisfiable.
2. If  $I$  is unsatisfiable then  $R(I)$  is not  $s$ -satisfiable.

We often say that a problem “has universal factor graphs” when it is UFG-NP-hard.

*Remark 2.13.* As a 3-Sat instance might have up to  $n^3$  clauses this definition is rather relaxed in how it handles size parameters. In this paper this does not matter as we do not keep track of degrees of various polynomials appearing in our proofs. If a more fine-grained theory was desired it would be useful to introduce also a parameter  $m$  for the number of clauses in the 3-Sat formula and see how this parameter enters into the size of the resulting factor graph family.

Our definitions here have minor technical differences with those in [FJ12]. This was done for reasons of presentation, and it is easy to check that all of the results in both this paper and [FJ12] hold for both (very similar) sets of definitions for universal factor graphs. In particular, it is not difficult to see that if  $\text{Max-CSP}(\Gamma)$  is  $(c, s)$ -UFG-NP-hard using a family  $\mathcal{F}$  then there is no  $\text{poly}(n)$ -size circuit family  $\{C_n\}$  that  $(c, s)$ -approximates  $\text{Max-CSP}(\Gamma)$  on instances with factor graph  $\mathcal{F}_n$  unless  $\text{NP} \subseteq \text{P/poly}$ .

To prove our hardness results, we start with some problem already known to have universal factor graphs, and then do a reduction to a CSP achieving the optimal approximation ratio. The key difference from the standard version of such reductions is that, given two instances with the same factor graph, our reductions must produce two instances with the same factor graphs.

**Definition 2.14.** A reduction  $R$  from  $\text{Max-CSP}(\Gamma)$  to  $\text{Max-CSP}(\Lambda)$  is *factor graph-preserving* if, whenever two  $\text{Max-CSP}(\Gamma)$  instances  $I$  and  $I'$  have the same factor graph, then  $R(I)$  and  $R(I')$  also have the same factor graph.

The key property of factor graph-preserving reductions is the following immediate fact.

**Fact 2.15.** *If  $\text{Max-CSP}(\Gamma)$  is  $(c, s)$ -UFG-NP-hard and there is a factor graph-preserving polynomial time reduction from  $(c, s)$ -approximating  $\text{Max-CSP}(\Gamma)$  to  $(c', s')$ -approximating  $\text{Max-CSP}(\Lambda)$ , then  $\text{Max-CSP}(\Lambda)$  is  $(c', s')$ -UFG-NP-hard.*

The starting point for our reductions is the following hardness result for Max-TSA.

**Theorem 2.16** ([Joz14]). *There is a constant  $s < 1$  such that Max-TSA is  $(1, s)$ -UFG-NP-hard.*

## 2.3 Analysis of Boolean Functions

We use standard notation, but in a slightly non-standard set-up. As mentioned in [Section 1.2](#) we are mostly concerned with analysing tables which are Boolean functions of Boolean functions, rather than functions taking generic bit strings as inputs. Of course there is no real difference between a Boolean function  $f \in \mathcal{F}_n$  and a bit string of length  $2^n$  as long as we identify  $\{0, 1\}^n$  and  $[2^n]$  but the notation looks slightly different. Furthermore, our choice to make Boolean functions take value in  $\{0, 1\}$  rather than  $\{-1, 1\}$  causes us to many times replace what would have been  $A(f)$  in the latter notation by  $(-1)^{A(f)}$  in our current notation. Let us turn to some definitions.

**Definition 2.17.** For  $\alpha \subseteq \{0, 1\}^n$  we have a character  $\chi_\alpha : \mathcal{F}_n \rightarrow \{-1, 1\}$  defined by

$$\chi_\alpha(f) = (-1)^{\sum_{\mathbf{x} \in \alpha} f(\mathbf{x})}.$$

**Definition 2.18.** For a Boolean function  $A : \mathcal{F}_n \rightarrow \{0, 1\}$  we define the Fourier coefficients by

$$\hat{A}_\alpha = \mathbb{E}_f [(-1)^{A(f)} \chi_\alpha(f)].$$

We have the Fourier inversion formula

$$(-1)^{A(f)} = \sum_{\alpha} \hat{A}_\alpha \chi_\alpha(f)$$

and Plancherel's identity  $\sum_{\alpha} \hat{A}_\alpha^2 = 1$ . The Boolean-valued function  $A(f)$  and the real-valued function  $(-1)^{A(f)}$  are of course just different views of the same mathematical object.

## 2.4 Parallel Repetition

An instance  $I = (X, C)$  of a Max-CSP can be naturally associated with a basic two-prover game. The verifier picks a random constraint  $(S, f) \in C$  and a uniformly random variable  $x_i$  from the tuple  $\mathbf{x}_S$ . It sends  $x_i$  to prover  $P_1$  and  $\mathbf{x}_S$  to prover  $P_2$ .  $P_1$  responds with a value for  $x_i$  and  $P_2$  responds with values for all the variables in  $\mathbf{x}_S$ . The verifier accepts if and only if the values given to  $x_i$  by the two provers are the same, and the value given to  $\mathbf{x}_S$  satisfies  $f$ .

If the instance  $I$  is satisfiable then the provers can win this game with probability 1 (perfect completeness), and if  $I$  is at most  $(1 - \delta)$ -satisfiable then they can win with probability at most  $1 - \delta/k$  where  $k$  is the arity of each constraint and thus we preserve soundness strictly smaller than one.

In the  $r$ -wise parallel repetition of this game the verifier chooses  $r$  random constraints which it sends to  $P_2$  and randomly one variable from each constraint which it sends to  $P_1$ . The provers respond with values for all the variables they are sent. The verifier accepts if and only if the values from  $P_2$  satisfy the constraints and match those sent by  $P_1$  on the common variables. The  $r$ -parallel repeated game has perfect completeness and soundness  $c^r$  for some  $c < 1$  [[Raz98](#)].

In our reductions we need a variant called *smooth* parallel repetition. In this version an extra set of  $t$  constraints are sent to both provers. Of course for this to be useful,  $P_2$  does not know which  $t$  of its  $t + r$  constraints are sent to  $P_1$ . The verifier now also checks that it gets the same values for the  $tk$  values requested from both provers. It is easy to see that these extra variables sent to both provers do not increase soundness which remains at most  $c^r$ .

We use smooth parallel repetition in order to ensure that for any two distinct answers,  $a_1$  and  $a_2$  sent by  $P_2$  it is unlikely that there is one answer by  $P_1$  that is accepting for both  $a_i$ . This is ensured by setting  $t$  significantly larger than  $r$ , while keeping both parameters constant independent of the number of variables in the CSP. Let us give a formal description.

**Definition 2.19.** Given a CSP instance  $I = (X, C)$  the  $(r, t)$ -smooth parallel repetition is the following two-prover game.

1. For  $j = 1, \dots, t + r$  choose a constraint  $c_j = (S_j, f_j) \in C$  uniformly at random.
2. For  $j = 1, \dots, r$  choose a variable  $x_{i_j} \in \mathbf{x}_{S_{t+j}}$  uniformly at random.
3. Send  $(\mathbf{x}_{S_j})_{j=1}^{t+r}$  in random order to  $P_2$ , and send both  $(\mathbf{x}_{S_j})_{j=1}^t$  and  $(x_{i_j})_{j=1}^r$  to  $P_1$ .
4. Receive values for the variables sent to each prover, and check that for each  $j$  the values  $\mathbf{a}_{S_j}$ , given by  $P_2$  to  $\mathbf{x}_{S_j}$  satisfy  $c_j$ , and that the two values given to each of  $(\mathbf{x}_{S_j})_{j=1}^t$  and  $(x_{i_j})_{j=1}^r$  by the two provers agree.

We denote the set of variables sent to  $P_1$  by  $U$  and the set sent to  $P_2$  by  $W$ .

The smoothness property of the repeated game is quantified by the following claim.

*Claim 2.20.* For a fixed set of variables  $W$  sent to  $P_2$  and any set of possible answers  $S \subseteq \{0, 1\}^W$  from  $P_2$  in the  $(r, t)$ -smooth parallel repetition, the probability (over the choice of  $U$  conditioned on  $W$ ) that there exists two different answers  $\mathbf{a}$  and  $\mathbf{a}'$  in  $S$  such that  $\mathbf{a}_U = \mathbf{a}'_U$  is at most  $\frac{|S|^2 r}{t+r}$ .

*Proof.* For any two fixed elements  $\mathbf{a}$  and  $\mathbf{a}'$  in  $S$ , the probability that  $\mathbf{a}_U = \mathbf{a}'_U$  is at most  $\frac{(k-1)r}{k(t+r)} \leq \frac{r}{t+r}$ . This follows as they differ in at least one coordinate and only  $(k-1)r$  of the  $k(t+r)$  coordinates are projected away. The claim follows by a union bound over all  $\binom{|S|}{2}$  pairs of elements of  $S$ .  $\square$

### 3 Functional Folding and Reduction Template

Most parts of our proofs are standard. We apply parallel repetition to a constraint-versus-variable two-prover proof and then code the answers of the provers by the long code. The main novelty is a new way of folding the long code in a factor graph-preserving way and we now describe this mechanism.

#### 3.1 Factor Graph-Preserving Folding of Long Codes

The full long code of a string  $\mathbf{x} \in \{0, 1\}^\ell$  is a table  $A : \mathcal{F}_\ell \rightarrow \{0, 1\}$  indexed by the set  $\mathcal{F}_\ell$  of all functions  $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ , where  $A(f)$  is supposed to take the value  $f(\mathbf{x})$ .

When giving a long code of a string  $\mathbf{x}$  that is supposed to satisfy some conditions, say  $h_i(\mathbf{x}) = b_i$  for  $i = 1, 2, \dots, r$ , it turns out to be essential to incorporate these conditions directly into the long code. The idea is to divide the functions into equivalence classes such that the value on one function in the equivalence class determines the value of all other functions in the same equivalence class. It is then sufficient to give only one bit for each equivalence class of functions. This has, in the past, been done in two different ways.

1. Put  $f$  and  $g$  in the same equivalence class if and only if  $f(\mathbf{x}) = g(\mathbf{x}) + \sum_{i=1}^r \sigma_i h_i(\mathbf{x})$  for some constants  $\sigma_i$ .
2. Put  $f$  and  $g$  in the same equivalence class if  $f(\mathbf{x}) = g(\mathbf{x})$  for all  $\mathbf{x}$  that satisfy  $h_i(\mathbf{x}) = b_i$  for all  $i$ . In many situations we also allow for the possibility that  $f(\mathbf{x}) = \neg g(\mathbf{x})$  for all such  $\mathbf{x}$ .

In particular, the former method, which we call linear folding, was used by Bellare et al. [BGS98] and the latter, which we call conditioning, by Håstad [Hås01]. The second method creates fewer equivalence classes.

In the current situation we want the construction to be factor graph-preserving and this turns out to be equivalent to the property that the equivalence classes do not depend on the unknown

negations. If each constraint comes from a CNF clause on some known set of variables  $S$  with unknown negations then it is easy to see that not even linear folding has this property. Once we write such a condition on the form  $h(\mathbf{x}_S) = b$  then the identity of  $h$  and hence the folding depends on the negations.

On the other hand if we start with an equational CSP, linear folding does have the desired property. Linear folding was sufficient in the case when the soundness analysis of the PCP protocol used established that a table given by the prover was close in Hamming distance to a correct long code. In most sharp inapproximability results one is not able to establish such a strong property and in such situations linear folding is not sufficient for the analysis of the protocol. However, even in the situation of equational CSPs it is easy to see that conditioning creates equivalence classes that are dependent on the right hand sides and thus it seems too strong to use in our setting.

The main new technical tool of this paper is to define an intermediate type of folding that is factor graph-preserving, but is still strong enough to be used instead of conditioning in many situations where a PCP is analyzed using Fourier analysis and/or the Invariance principle. We proceed to define this folding which is based on the following equivalence relation.

**Definition 3.1.** Let  $H = \{h_1, \dots, h_r\} \subseteq \mathcal{F}_\ell$ . Two functions  $f, g \in \mathcal{F}_\ell$  are  $H$ -equivalent if there exists  $F : \{0, 1\}^r \rightarrow \{0, 1\}$  such that

$$f(\mathbf{x}) = g(\mathbf{x}) + F(h_1(\mathbf{x}), \dots, h_r(\mathbf{x}))$$

for all  $\mathbf{x} \in \{0, 1\}^\ell$ .

It is easy to check that this definition gives an equivalence relation, and we refer to the classes of this relation as  $H$ -equivalence classes. Now we can define the folding of  $A$  over all functions on the constraints.

**Definition 3.2.** Let  $A : \mathcal{F}_\ell \rightarrow \{0, 1\}$  be a supposed long code,  $\{h_i(\mathbf{x}) = b_i\}_{i=1}^r$  be a set of constraints,  $\mathbf{b} = (b_1, \dots, b_r) \in \{0, 1\}^r$ , and  $H = \{h_1, \dots, h_r\}$ . For each class of  $H$ -equivalent functions we choose one representative. The folding of  $A$  over all functions on  $H$  with respect to  $\mathbf{b}$ , denoted by  $A_{H,\mathbf{b}}$ , is now defined as follows. If  $g$  is the representative for an  $H$ -equivalence class containing some function  $f$ , then we define

$$A_{H,\mathbf{b}}(f) \stackrel{\text{def}}{=} A(g) + F(b_1, b_2, \dots, b_r).$$

where  $f = g + F(h_1, h_2, \dots, h_r)$ .

We call this folding *functional folding*, and it is easy to see that if  $A_{H,\mathbf{b}}$  is functionally folded, then changing the choice of representative for each  $H$  equivalence class does not change the function. Note that even if  $r = 0$  we always fold over the function that is identically true.

Given  $A$ , we simulate queries to  $A_{H,\mathbf{b}}$  in the standard way. Whenever we want to read the value of  $A_{H,\mathbf{b}}(f)$ , we find the representative  $g = f + F(h_1, \dots, h_r)$  of the  $H$ -equivalence class for  $f$ , and return the value  $A(g) + F(b_1, b_2, \dots, b_r)$ . Since the equivalence classes depend only on the  $h_i$  and not on the  $b_i$ , we immediately have that this folding is factor graph-preserving when we start with an equational CSP.

**Fact 3.3.** *Let  $A, H$  and  $\mathbf{b}$  be as above. If a PCP verifier simulates queries to  $A_{H,\mathbf{b}}$  by querying the table  $A$ , then the query locations depend only on  $H$ .*

While this folding is factor graph-preserving, it is not yet clear that it is actually useful in enforcing the constraints. The next lemma shows that the folding does indeed enforce the constraints in the sense that all non-zero Fourier coefficients correspond to sets of assignments with an odd number of them satisfying the constraints.

**Lemma 3.4.** *Let  $A : \mathcal{F}_\ell \rightarrow \{0, 1\}$  be a supposed long code,  $\{h_i(\mathbf{x}) = b_i\}_{i=1}^r$  be a set of constraints,  $\mathbf{b} = (b_1, \dots, b_r) \in \{0, 1\}^r$ , and  $H = \{h_1, \dots, h_r\}$ . Let  $A_{H,\mathbf{b}}$  be  $A$  folded over all functions on  $H$  with respect to  $\mathbf{b}$ . If  $\hat{A}_{H,\mathbf{b}}(\alpha) \neq 0$  then the number of  $\mathbf{x} \in \alpha$  that simultaneously satisfy  $h_i(\mathbf{x}) = b_i$  for all  $i$  is odd.*

*Proof.* Let  $h(\mathbf{x}) = (h_1(\mathbf{x}), \dots, h_r(\mathbf{x}))$ . Suppose that the number of  $\mathbf{x} \in \alpha$  that satisfy all of the constraints  $h_i(\mathbf{x}) = b_i$  is even. Recall that

$$\hat{A}_{H,b}(\alpha) = \mathbb{E}_{f \in \mathcal{F}_\ell} \left[ (-1)^{A_{H,b}(f)} \chi_\alpha(f) \right] = \mathbb{E}_{f \in \mathcal{F}_\ell} \left[ (-1)^{A_{H,b}(f)} \prod_{\mathbf{x} \in \alpha} (-1)^{f(\mathbf{x})} \right]. \quad (3.1)$$

Consider the function  $\mathbf{1}_b : \{0, 1\}^r \rightarrow \{0, 1\}$  which is one only at the point  $\mathbf{b}$ . Now we pair up the functions  $f$  and  $f + \mathbf{1}_b(h)$  in (3.1). By the folding we have

$$A_{H,b}(f + \mathbf{1}_b(h)) = A_{H,b}(f) + \mathbf{1}_b(\mathbf{b}) = A_{H,b}(f) + 1.$$

In particular this implies that  $(-1)^{A_{H,b}(f + \mathbf{1}_b(h))} = -((-1)^{A_{H,b}(f)})$ . On the other hand, since the number of  $\mathbf{x} \in \alpha$  satisfying  $h(\mathbf{x}) = \mathbf{b}$  is even we have

$$\prod_{\mathbf{x} \in \alpha} (-1)^{f(\mathbf{x}) + \mathbf{1}_b(h(\mathbf{x}))} = \prod_{\mathbf{x} \in \alpha} (-1)^{f(\mathbf{x})} (-1)^{\mathbf{1}_b(h(\mathbf{x}))} = \prod_{\mathbf{x} \in \alpha} (-1)^{f(\mathbf{x})}$$

Thus the terms corresponding to  $f$  and  $f + \mathbf{1}_b(h)$  cancel in pairs and the expectation (3.1) is zero.  $\square$

**Remark.** By a similar argument, the number of elements in  $\beta$  that give any other fixed right hand sides is even. As this is not needed in our arguments we leave the verification of this to the reader.

Lemma 3.4 shows that if we use functional folding then for any non-zero Fourier coefficient at least one element satisfies all the constraints. This should be compared with conditioning where *all* elements in a non-zero Fourier coefficient satisfies all the constraints. On the other hand with linear folding where we have no such guarantees and it is difficult to find an assignment that satisfies the constraints given a non-zero Fourier coefficient.

## 3.2 Basic Setup of Hardness Reductions

It turns out that by using some minor modifications, the guarantee of Lemma 3.4 is sufficient to analyze many of the classical protocols.

We follow the standard setup: start with a Max-CSP having some hardness of approximation, apply parallel repetition to create a two-prover game with perfect completeness and arbitrarily small soundness, and then long code the answers of the two provers and test the long codes using constraints from the target CSP. In order to apply functional folding to the long codes, we start with an equational CSP and in particular use the hardness of the Max-TSA problem (Theorem 2.16), and in order to be able to use the weaker guarantee of functional folding as compared to conditioning we use smooth parallel repetition (Definition 2.19).

As all the reductions follow the same initial steps and only the exact choice of constraints varies from CSP to CSP, we summarize these first steps and the notation used in the following definition.

**Definition 3.5** (Reduction template for Max-CSPs with universal factor graphs). Given a Max-TSA instance  $I = (X, C)$  and parameters  $r$  and  $t$ , we construct a new set  $X'$  of Boolean variables as follows.

First, form the  $(r, t)$ -smooth parallel repeated game of  $I$  as in Definition 2.19. For any set of variables  $U \subseteq X$  that may be sent to  $P_1$ , we introduce a supposed long code  $A_U : \mathcal{F}_U \rightarrow \{0, 1\}$  of the assignment to  $U$ . We write  $A_U^{\boxplus}$  for  $A_U$  functionally folded over all constraints of  $I$  induced by  $U$ . For a set  $W \subseteq X$  that may be sent to  $P_2$  we define  $B_W$  and  $B_W^{\boxplus}$  analogously.

The variables  $X'$  consist of all values of  $A_U(f)$  and  $B_W(g)$  over all  $U, W, f \in \mathcal{F}_U$  and  $g \in \mathcal{F}_W$ .

In the concrete reductions in the rest of the paper, a set  $U$  (resp.  $W$ ) always refers to a query to  $P_1$  (resp.  $P_2$ ) in the repeated game.

Note that, as per the definition of functional folding in the previous section, queries to  $A_U^{\boxplus}$  and  $B_W^{\boxplus}$  are made by querying fixed entries of  $A_U$  and  $B_W$  (depending only on the factor graph of  $I$ ) and then possibly negating the result (depending on the right hand sides of  $I$ ).

## 4 Classical Optimal Inapproximability

Using the folding introduced in the previous section we prove classical optimal inapproximability results with universal factor graphs. The first two problems, Max-3-Sat and Max-3-Lin, were first proven to be approximation resistant in [Hås01]. The third problem, Max-TSA, did not appear in that original paper, but the techniques used follow the standard dictatorship testing analysis via the Fourier transform introduced there. In each case, we start with the reduction template from Definition 3.5 and then construct a PCP verifier for the target problem.

In the hardness results for Max-3-Sat and Max-TSA, we use the notation  $\beta_{\oplus U}$  to denote the set of vectors  $\mathbf{x} \in \beta_U$  such that there is an odd number of  $\mathbf{y} \in \beta$  with  $\mathbf{y}_U = \mathbf{x}$ . In the notation of [Hås01], this is the  $\pi_2$  operator.

### 4.1 Max-3-Lin

**Theorem 4.1.** *For any  $\varepsilon > 0$ , Max-3-Lin is  $(1 - \varepsilon, \frac{1}{2} + \varepsilon)$ -UFG-NP-hard.*

The long code test we use is exactly the same as the original test of Håstad in [Hås01], but as explained in Section 3, we can only access the functional foldings of the long codes, and not condition them on the constraints of the underlying CSP.

**Definition 4.2.** The Max-3-Lin PCP verifier does the following:

1. Pick a random pair of sets  $(U, W)$  sent to the two provers in the parallel game.
2. Pick a uniform random function  $f \in \mathcal{F}_U$  and a uniform random  $g_1 \in \mathcal{F}_W$ .
3. Define  $g_2 \in \mathcal{F}_W$  by setting  $g_2(\mathbf{y}) = g_1(\mathbf{y}) + f(\mathbf{y}_U)$  with probability  $1 - \varepsilon$  and the negation of this value otherwise, independently for each  $\mathbf{y} \in \{0, 1\}^W$ .
4. Accept if and only if  $A_U^{\boxplus}(f) + B_W^{\boxplus}(g_1) + B_W^{\boxplus}(g_2) = 0$ .

The tests defined above correspond to checking three variable linear equations over  $\mathbb{F}_2$ , and so the PCP defines a Max-3-Lin instance. Also we have immediately by Fact 3.3 that the overall reduction from Max-TSA to Max-3-Lin given by this PCP is factor graph-preserving.

Now we analyze the completeness and soundness of the PCP.

**Lemma 4.3.** *The completeness of the Max-3-Lin verifier in Definition 4.2 is  $1 - \varepsilon$ .*

*Proof.* Let  $\mathbf{a}$  be a satisfying assignment to the original Max-TSA constraints. Then we set  $A_U(f) = f(\mathbf{a}_U)$  and  $B_W(g) = f(\mathbf{a}_W)$  for all subsets of variables in the two prover game. Note that  $B_W(g) = B_W^{\boxplus}(g)$  since functional folding does not affect a true long code of a satisfying assignment. Second, with probability  $1 - \varepsilon$  we have  $A_U^{\boxplus}(f) + B_W^{\boxplus}(g_1) + A_W^{\boxplus}(g_2) = f(\mathbf{a}_U) + g_1(\mathbf{a}_W) + g_2(\mathbf{a}_W) = 0$  and hence the test accepts with probability  $1 - \varepsilon$ .  $\square$

We now prove soundness for the test.

**Lemma 4.4.** *Let  $\varepsilon, \delta > 0$  and set  $C_{\varepsilon, \delta} \stackrel{\text{def}}{=} (2\varepsilon)^{-1} \log(\frac{2}{\delta})$ . For  $r > 0$  set  $t = r \cdot \frac{4}{\delta^2} C_{\varepsilon, \delta}^2$ . If the Max-3-Lin PCP verifier accepts with probability at least  $\frac{1+\delta}{2}$  then there is a strategy for the provers in the  $(r, t)$ -smooth parallel repeated game that causes the verifier to accept with probability at least  $\frac{\delta^4}{16C_{\varepsilon, \delta}}$ .*

*Proof.* For notational convenience, throughout this proof we write  $A_U$  and  $B_W$  instead of  $A_U^{\boxplus}$  and  $B_W^{\boxplus}$ , but it is important to keep in mind that all long codes are folded. Now the standard analysis gives that if the test accepts with probability  $(1 + \delta)/2$  then

$$\delta = \mathbb{E}_{U, W} \left[ \mathbb{E}_{f, g_1, g_2} \left[ (-1)^{A_U(f)} (-1)^{B_W(g_1)} (-1)^{B_W(g_2)} \right] \right]$$

$$\begin{aligned}
&= \mathbb{E}_{U,W} \left[ \sum_{\alpha, \beta_1, \beta_2} \hat{A}_U(\alpha) \hat{B}_W(\beta_1) \hat{B}_W(\beta_2) \mathbb{E}_{f, g_1, g_2} [\chi_\alpha(f) \chi_{\beta_1}(g_1) \chi_{\beta_2}(g_2)] \right] \\
&= \mathbb{E}_{U,W} \left[ \sum_{\beta} \hat{A}_U(\beta_{\oplus U}) \hat{B}_W(\beta)^2 \mathbb{E} [\chi_{\beta_{\oplus U}}(f) \chi_{\beta}(g_1 + g_2)] \right] \\
&= \mathbb{E}_{U,W} \left[ \sum_{\beta} \hat{A}_U(\beta_{\oplus U}) \hat{B}_W(\beta)^2 (1 - 2\varepsilon)^{|\beta|} \right]
\end{aligned}$$

and an application of Cauchy-Schwartz and Plancherel gives

$$\delta^2 \leq \mathbb{E}_{U,W} \left[ \sum_{\beta} \hat{A}_U(\beta_{\oplus U})^2 \hat{B}_W(\beta)^2 (1 - 2\varepsilon)^{2|\beta|} \right]. \quad (4.1)$$

Note that the contribution to the inner sum from  $\beta$  that are larger than  $C_{\varepsilon, \delta}$  is at most  $\delta^2/4$ . Thus we have

$$\frac{3}{4} \delta^2 \leq \mathbb{E}_{U,W} \left[ \sum_{|\beta| \leq C_{\varepsilon, \delta}} \hat{A}_U(\beta_{\oplus U})^2 \hat{B}_W(\beta)^2 \right]. \quad (4.2)$$

The strategies for the two provers are as follows. With probability  $\hat{A}_U(\alpha)^2$  prover  $P_1$  outputs a random  $\mathbf{x} \in \alpha$ , and with probability  $\hat{B}_W(\beta)^2$  prover  $P_2$  outputs a random  $\mathbf{y} \in \beta$  which satisfies all the constraints on  $W$ . By [Lemma 3.4](#) for any  $\beta$  with  $\hat{B}_W(\beta) \neq 0$  there are an odd number of assignments  $\mathbf{y} \in \beta$  satisfying the constraints on  $W$ . In particular, there is at least one assignment  $\mathbf{y}^* \in \beta$  satisfying all the constraints. However, if  $\mathbf{y}^*$  and another assignment  $\mathbf{y} \in \beta$  collide under the map  $\mathbf{y} \rightarrow \mathbf{y}_U$ , then there may be no corresponding assignment  $\mathbf{x} = \mathbf{y}_U^*$  in  $\beta_{\oplus U}$ .

We now analyze the contribution to [\(4.2\)](#) of terms where such collisions occur. By [Claim 2.20](#), we have that for each fixed  $W$  and  $\beta$ , the probability over the choice of  $U$  of a collision is at most  $|\beta|^2 \frac{r}{t+r}$ . Let  $S_{W, \beta}$  be the set of  $U$  which cause a collision. Then

$$\begin{aligned}
\sum_{|\beta| \leq C_{\varepsilon, \delta}} \mathbb{E}_U \left[ \mathbf{1}(U \in S_{W, \beta}) \hat{A}_U(\beta_{\oplus U})^2 \hat{B}_W(\beta)^2 \right] &\leq \sum_{|\beta| \leq C_{\varepsilon, \delta}} \mathbb{E}_U \left[ \mathbf{1}(U \in S_{W, \beta}) \right] \\
&\leq C_{\varepsilon, \delta}^2 \frac{r}{t+r} \leq \frac{\delta^2}{4}.
\end{aligned}$$

Combining this with [\(4.2\)](#) yields

$$\frac{\delta^2}{2} \leq \mathbb{E}_{U,W} \left[ \sum_{|\beta| \leq C_{\varepsilon, \delta}} \mathbf{1}(U \notin S_{W, \beta}) \hat{A}_U(\beta_{\oplus U})^2 \hat{B}_W(\beta)^2 \right].$$

Since the inner sum is non-negative and bounded by one, we have by Markov's inequality that with probability at least  $\frac{\delta^2}{4}$  over the choice of  $U$  and  $W$

$$\frac{\delta^2}{4} \leq \sum_{|\beta| \leq C_{\varepsilon, \delta}} \mathbf{1}(U \notin S_{W, \beta}) \hat{A}_U(\beta_{\oplus U})^2 \hat{B}_W(\beta)^2.$$

This implies that for each such pair  $(U, W)$ , with probability at least  $\frac{\delta^2}{4}$  the following events all occur:

- $P_2$  chooses some  $\beta$  with  $|\beta| \leq C_{\varepsilon, \delta}$ .
- $P_1$  chooses  $\alpha = \beta_{\oplus U}$ .
- $U \notin S_{W, \beta}$  i.e. no two elements of  $\beta$  collide under projection to  $U$ .

If no two elements of  $\beta$  collide under the map  $\mathbf{y} \mapsto \mathbf{y}_U$ , then the projection  $\mathbf{y}_U^*$  of the good assignment  $\mathbf{y}^* \in \beta$  satisfying all the constraints on  $W$  is also an element of  $\beta_{\oplus U}$ . Therefore  $P_1$  chooses  $\mathbf{y}_U^*$  with probability at least  $|\alpha|^{-1} \geq |\beta|^{-1} \geq C_{\varepsilon, \delta}^{-1}$ . When this happens, the responses of  $P_1$  and  $P_2$  to  $U$  and  $W$  are accepted in the parallel repeated game. Thus, as these three events happen for a  $\frac{\delta^2}{4}$  fraction of all queries  $(U, W)$ , the strategy of the two provers is accepted with probability at least  $\frac{\delta^4}{16} C_{\varepsilon, \delta}^{-1}$ , as desired.  $\square$

The reduction can now be completed by recalling that the parallel repeated game has soundness  $c^r$  for some constant  $c$ . But by [Lemma 4.4](#) (setting  $\delta = 2\varepsilon$ ) the value of the repeated game is larger than  $c^r$  for sufficiently large  $r$ , unless fewer than a  $\frac{1}{2} + \varepsilon$  fraction of equations in the Max-3-Lin instance can be satisfied.

## 4.2 Max-3-Sat

As with Max-3-Lin, we can use our new folding along with a variant of the Max-3-Sat test in [\[Hås01\]](#) to obtain UFG-NP-hardness. Here we do not follow exactly the same long code test as [\[Hås01\]](#) but instead follow a subsequent simplified test and analysis using smooth label cover [\[Kho02, GK05\]](#).

**Theorem 4.5.** *For any  $\varepsilon > 0$ , Max-3-Sat is  $(1, \frac{7}{8} + \varepsilon)$ -UFG-NP-hard.*

The PCP we construct for Max-3-Sat is described below.

**Definition 4.6.** The Max-3-Sat PCP verifier does the following:

1. Pick a random pair of sets  $(U, W)$  sent to the two provers in the parallel game.
2. Pick a uniform random function  $f \in \mathcal{F}_U$  and a uniform random  $g_1 \in \mathcal{F}_W$ .
3. Pick a function  $g_2 \in \mathcal{W}$  as follows. For each  $\mathbf{y} \in \{0, 1\}^W$ , if  $f(\mathbf{y}_U) = 0$  then assign  $g_2(\mathbf{y}) = g_1(\mathbf{y}) + 1$ . If  $f(\mathbf{y}_U) = 1$  then with probability  $1 - \varepsilon$  assign  $g_2(\mathbf{y}) = g_1(\mathbf{y})$  and otherwise assign  $g_2(\mathbf{y}) = g_1(\mathbf{y}) + 1$ .
4. Accept unless  $A_U^{\boxplus}(f) = B_W^{\boxplus}(g_1) = B_W^{\boxplus}(g_2) = 0$ .

It is easy to prove that this test has perfect completeness, and as the proof is near-identical to that of [Lemma 4.3](#) we omit it.

**Lemma 4.7.** *The completeness of the Max-3-Sat verifier in [Definition 4.6](#) is 1.*

**Lemma 4.8.** *Let  $\delta > 0$ , and let  $0 < \varepsilon < (\frac{\delta}{32})^2 \log(\frac{16}{\delta})^{-2}$ . Set  $C_{\varepsilon, \delta} = \frac{1}{\varepsilon} \log(\frac{16}{\delta})$ . For  $r > 0$ , let  $t = r \cdot \frac{16}{\delta} C_{\varepsilon, \delta}^2$ . If the Max-3-Sat PCP verifier accepts with probability at least  $\frac{7+\delta}{8}$  then there is a strategy for the provers in the  $(r, t)$ -smooth parallel repeated game that causes the verifier to accept with probability at least  $\frac{\delta^4}{64C_{\varepsilon, \delta}}$ .*

*Proof.* As in the soundness analysis for Max-3-Lin, we write  $A_U$  and  $B_W$  instead of  $A_U^{\boxplus}$  and  $B_W^{\boxplus}$  throughout this proof in order to keep the notation manageable. The probability that the test accepts is given by

$$\mathbb{E}_{U, W} \left[ \mathbb{E}_{f, g_1, g_2} \left[ 1 - \frac{1}{8} (1 + (-1)^{A_U(f)}) (1 + (-1)^{B_W(g_1)}) (1 + (-1)^{B_W(g_2)}) \right] \right]. \quad (4.3)$$

Recall that the long codes  $A_U$  and  $B_W$  are folded over true, so  $\mathbb{E} [(-1)^{A_U(f)}] = \mathbb{E} [(-1)^{B_W(g_i)}] = 0$ . Furthermore since  $f$  and  $g_i$  are independent we have  $\mathbb{E} [(-1)^{A_U(f)}(-1)^{B_W(g_i)}] = \mathbb{E} [(-1)^{A_U(f)}] \mathbb{E} [(-1)^{B_W(g_i)}] = 0$ . In other words, after expanding (4.3), any non-constant term which does not contain  $(-1)^{B_W(g_1)}(-1)^{B_W(g_2)}$  becomes 0, so the acceptance probability equals

$$\frac{7}{8} - \frac{1}{8} \mathbb{E}_{U,W} \left[ \mathbb{E}_{f,g_1,g_2} \left[ (-1)^{B_W(g_1)}(-1)^{B_W(g_2)} + (-1)^{A_U(f)}(-1)^{B_W(g_1)}(-1)^{B_W(g_2)} \right] \right] \geq \frac{7+\delta}{8} \quad (4.4)$$

We analyze the two terms in the above expectation one at a time.

*Claim 4.9.*

$$\left| \mathbb{E}_{U,W} \left[ \mathbb{E}_{g_1,g_2} \left[ (-1)^{B_W(g_1)}(-1)^{B_W(g_2)} \right] \right] \right| \leq \delta/4. \quad (4.5)$$

*Proof.* For a fixed choice of  $U$  and  $W$ , writing out the Fourier expansion of the inner expectation of this term yields:

$$\begin{aligned} \mathbb{E}_{g_1,g_2} \left[ (-1)^{B_W(g_1)}(-1)^{B_W(g_2)} \right] &= \sum_{\beta_1,\beta_2} \hat{B}_W(\beta_1)\hat{B}_W(\beta_2) \mathbb{E} [\chi_{\beta_1}(g_1)\chi_{\beta_2}(g_2)] \\ &= \sum_{\beta} \hat{B}_W(\beta)^2 \mathbb{E} [\chi_{\beta}(g_1 + g_2)], \end{aligned} \quad (4.6)$$

where we used the fact that for  $\beta_1 \neq \beta_2$  the expectation inside the above sum is zero. Now observe that if  $\mathbf{y}_U \neq \mathbf{y}'_U$ , then  $g_1(\mathbf{y}) + g_2(\mathbf{y})$  is independent of  $g_1(\mathbf{y}') + g_2(\mathbf{y}')$ . Using this fact we have

$$\mathbb{E} [\chi_{\beta}(g_1 + g_2)] = \mathbb{E} \left[ \prod_{\mathbf{y} \in \beta} (-1)^{g_1(\mathbf{y})+g_2(\mathbf{y})} \right] = \prod_{\mathbf{x} \in \beta_U} \mathbb{E} \left[ \prod_{\mathbf{y} \in \beta: \mathbf{y}_U = \mathbf{x}} (-1)^{g_1(\mathbf{y})+g_2(\mathbf{y})} \right]$$

Let  $s_{\mathbf{x}}$  be the number of  $\mathbf{y} \in \beta$  with  $\mathbf{y}_U = \mathbf{x}$ . The expectation over  $\mathbf{y}$  which project to the same  $\mathbf{x}$  is equal to

$$\mathbb{E} \left[ \prod_{\mathbf{y} \in \beta: \mathbf{y}_U = \mathbf{x}} (-1)^{g_1(\mathbf{y})+g_2(\mathbf{y})} \right] = \frac{1}{2}((-1)^{s_{\mathbf{x}}} + (1 - 2\varepsilon)^{s_{\mathbf{x}}}).$$

We bound such terms in two different ways depending on the size of  $\beta$ . First, observe that  $|\frac{1}{2}((-1)^{s_{\mathbf{x}}} + (1 - 2\varepsilon)^{s_{\mathbf{x}}})| \leq 1 - \varepsilon$  and so  $|\mathbb{E} [\chi_{\beta}(g_1 + g_2)]| \leq (1 - \varepsilon)^{|\beta_U|}$ . Applying Claim 2.20 to the first  $C_{\varepsilon,\delta}$  elements of  $\beta$  for  $|\beta| > C_{\varepsilon,\delta}$ , we have  $|\beta_U| \geq C_{\varepsilon,\delta}$  except with probability at most  $\frac{r}{t+r}C_{\varepsilon,\delta}^2 \leq \delta/16$  over the choice of  $U$ . Thus, averaging over  $U$  and  $W$  we get

$$\begin{aligned} \left| \mathbb{E}_{U,W} \left[ \sum_{\beta: |\beta| > C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 \mathbb{E} [\chi_{\beta}(g_1 + g_2)] \right] \right| &\leq \mathbb{E}_{U,W} \left[ \sum_{\beta: |\beta| > C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 (1 - \varepsilon)^{|\beta_U|} \right] \\ &\leq \mathbb{E}_W \left[ \sum_{\beta: |\beta| > C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 ((1 - \varepsilon)^{C_{\varepsilon,\delta}} + \delta/16) \right] \leq \frac{\delta}{8}, \end{aligned} \quad (4.7)$$

where the last inequality used our choice of  $C_{\varepsilon,\delta}$  and Plancherel. We turn to  $\beta$  of size at most  $C_{\varepsilon,\delta}$  and, again by Claim 2.20, except with probability at most  $\frac{r}{t+r}|\beta|^2 \leq \delta/16$  over the choice of  $U$ , there is at least one  $\mathbf{y}$  in  $\beta$  that does not collide with any other  $\mathbf{y}' \in \beta$ . Letting  $\mathbf{x} = \mathbf{y}_U$  we then have

$s_{\mathbf{x}} = 1$ . For such “good” choices of  $U$  we have  $\frac{1}{2}((-1)^{s_{\mathbf{x}}} + (1 - 2\varepsilon)^{s_{\mathbf{x}}}) = -\varepsilon$ . Thus after averaging over  $U$  and  $W$  we have that

$$\left| \mathbb{E}_{U,W} \left[ \sum_{\beta:|\beta|\leq C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 \mathbb{E} [\chi_{\beta}(g_1 + g_2)] \right] \right| \leq \varepsilon + \delta/16 < \frac{\delta}{8}. \quad (4.8)$$

Adding up (4.7) and (4.8) we obtain (4.5) (via (4.6)).  $\square$

We now analyze the second term.

*Claim 4.10.*

$$\left| \mathbb{E}_{U,W} \left[ \mathbb{E}_{g_1,g_2} \left[ (-1)^{A_U(f)} (-1)^{B_W(g_1)} (-1)^{B_W(g_2)} \right] \right] \right| \leq \frac{\delta}{4} + \mathbb{E}_{U,W} \left[ \sum_{|\beta|\leq C_{\varepsilon,\delta}} \hat{A}_U(\beta_U)^2 \hat{B}_W(\beta)^2 (1 - \varepsilon)^{2|\beta_U|} \right]^{1/2} \quad (4.9)$$

*Proof.* For a fixed choice of  $U$  and  $W$ , writing the Fourier expansion of the inner expectation yields

$$\begin{aligned} \mathbb{E}_{g_1,g_2} \left[ (-1)^{A_U(f)} (-1)^{B_W(g_1)} (-1)^{B_W(g_2)} \right] &= \sum_{\alpha,\beta_1,\beta_2} \hat{A}_U(\alpha) \hat{B}_W(\beta_1) \hat{B}_W(\beta_2) \mathbb{E} [\chi_{\alpha}(f) \chi_{\beta_1}(g_1) \chi_{\beta_2}(g_2)] \\ &= \sum_{\beta, \alpha \subseteq \beta_U} \hat{A}_U(\alpha) \hat{B}_W(\beta)^2 \mathbb{E} [\chi_{\alpha}(f) \chi_{\beta}(g_1 + g_2)] \end{aligned} \quad (4.10)$$

where we have used the fact that the expectation is zero unless  $\beta_1 = \beta_2 = \beta$  and  $\alpha \subseteq \beta_U$ . Let us define  $E(\alpha, \beta) \stackrel{\text{def}}{=} \mathbb{E} [\chi_{\alpha}(f) \chi_{\beta}(g_1 + g_2)]$  to denote the inner expectation above (and note that this function depends on  $U$  and  $W$ ). Next let  $s_{\mathbf{x}}$  be the number of  $\mathbf{y} \in \beta$  such that  $\mathbf{y}_U = \mathbf{x}$ . Then for  $\alpha \subseteq \beta_U$

$$E(\alpha, \beta) = \prod_{\mathbf{x} \in \alpha} \frac{1}{2}((-1)^{s_{\mathbf{x}}} - (1 - 2\varepsilon)^{s_{\mathbf{x}}}) \prod_{\mathbf{x} \in \beta_U \setminus \alpha} \frac{1}{2}((-1)^{s_{\mathbf{x}}} + (1 - 2\varepsilon)^{s_{\mathbf{x}}}). \quad (4.11)$$

Observe that this implies  $|E(\alpha, \beta)| \leq (1 - \varepsilon)^{|\beta_U|}$ , since every factor in the product is bounded in magnitude by  $1 - \varepsilon$ . Further note that

$$\begin{aligned} \sum_{\alpha \subseteq \beta_U} E(\alpha, \beta)^2 &= \prod_{\mathbf{x} \in \beta_U} \left( \left( \frac{1}{2}((-1)^{s_{\mathbf{x}}} - (1 - 2\varepsilon)^{s_{\mathbf{x}}}) \right)^2 + \left( \frac{1}{2}((-1)^{s_{\mathbf{x}}} + (1 - 2\varepsilon)^{s_{\mathbf{x}}}) \right)^2 \right) \\ &\leq (1 - \varepsilon)^{|\beta_U|} \end{aligned}$$

where the final inequality follows from the fact that each factor above has the form  $a^2 + b^2$  where both  $|a|$  and  $|b|$  are bounded by  $1 - \varepsilon$ , and  $|a| + |b| = 1$ . Therefore each factor is bounded by  $(1 - \varepsilon)^2 + \varepsilon^2$  which is at most  $1 - \varepsilon$  whenever  $\varepsilon \leq 1/2$ .

As in [Claim 4.9](#) we split the sum depending on the size of  $\beta$ . First by Cauchy-Schwarz and Plancherel we have

$$\begin{aligned} \sum_{\beta:|\beta|\geq C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 \sum_{\alpha \subseteq \beta_U} \hat{A}_U(\alpha) E(\alpha, \beta) &\leq \sum_{\beta:|\beta|\geq C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 \left( \sum_{\alpha \subseteq \beta_U} \hat{A}_U(\alpha)^2 \right)^{1/2} \left( \sum_{\alpha \subseteq \beta_U} E(\alpha, \beta)^2 \right)^{1/2} \\ &\leq \sum_{\beta:|\beta|\geq C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 (1 - \varepsilon)^{|\beta_U|} \end{aligned} \quad (4.12)$$

From (4.7) it follows that, when averaging over  $U$  and  $W$ , (4.12) is bounded in absolute value by  $\delta/8$ , i.e.,

$$\left| \mathbb{E}_{U,W} \left[ \sum_{\beta:|\beta|\geq C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 \sum_{\alpha\subseteq\beta_U} \hat{A}_U(\alpha)E(\alpha,\beta) \right] \right| \leq \delta/8. \quad (4.13)$$

Turning to the low-degree terms where  $|\beta| \leq C_{\varepsilon,\delta}$ , we apply Claim 2.20 to conclude that except with probability  $\frac{r}{t+r}|\beta| \leq \delta/16$  over the choice of  $U$ , there are no collisions between any  $\mathbf{y}, \mathbf{y}'$  in  $\beta$  under the map  $\mathbf{y} \mapsto \mathbf{y}_U$ . Letting  $\mathbf{x} = \mathbf{y}_U$  we then have  $s_{\mathbf{x}} = 1$ . So restricting to such good choices of  $U$  we have  $\frac{1}{2}((-1)^{s_{\mathbf{x}}} + (1-2\varepsilon)^{s_{\mathbf{x}}}) = -\varepsilon$ . Since all other factors in the product in (4.11) are bounded in magnitude by 1 we have for such “good”  $U$  that  $|E(\alpha, \beta)| \leq \varepsilon^{|\beta_U \setminus \alpha|}$ . Thus, in this case summing over all  $\alpha \subsetneq \beta_U$  yields

$$\begin{aligned} \sum_{|\beta|\leq C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 \sum_{\alpha\subsetneq\beta_U} \hat{A}_U(\alpha)E(\alpha,\beta) &\leq \sum_{|\beta|\leq C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 \left( \sum_{\alpha\subsetneq\beta_U} \hat{A}_U(\alpha)^2 \right)^{1/2} \left( \sum_{\alpha\subsetneq\beta_U} E(\alpha,\beta)^2 \right)^{1/2} \\ &\leq \sum_{|\beta|\leq C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 \left( \sum_{\alpha\subsetneq\beta_U} \varepsilon^{2|\beta_U \setminus \alpha|} \right)^{1/2} \\ &= \sum_{|\beta|\leq C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 \left( (1+\varepsilon^2)^{|\beta_U|} - 1 \right)^{1/2} \\ &\leq \sum_{|\beta|\leq C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 2\varepsilon C_{\varepsilon,\delta}^{1/2}, \end{aligned}$$

since  $(1+a)^b \leq 1+2ab$  whenever  $0 \leq ab \leq \frac{1}{32}$ . For the remaining at most  $\delta/16$  fraction of choices of  $U$  where collisions occur, we simply use the previous bound  $\sum_{\alpha\subseteq\beta_U} E(\alpha,\beta)^2 \leq (1-\varepsilon)^{|\beta_U|} \leq 1$ . Thus averaging over  $U$  and  $W$  yields

$$\left| \mathbb{E}_{U,W} \left[ \sum_{|\beta|\leq C_{\varepsilon,\delta}} \sum_{\alpha\subsetneq\beta_U} \hat{A}_U(\alpha) \hat{B}_W(\beta)^2 E(\alpha,\beta) \right] \right| \leq \mathbb{E}_W \left[ \sum_{|\beta|\leq C_{\varepsilon,\delta}} \hat{B}_W(\beta)^2 \left( 2\varepsilon C_{\varepsilon,\delta}^{1/2} + \delta/16 \right) \right] \leq \frac{\delta}{8}. \quad (4.14)$$

Finally, applying Cauchy-Schwarz and Plancherel to the sum where  $\alpha = \beta_U$  yields

$$\left| \mathbb{E}_{U,W} \left[ \sum_{|\beta|\leq C_{\varepsilon,\delta}} \hat{A}_U(\beta_U) \hat{B}_W(\beta)^2 E(\alpha,\beta) \right] \right| \leq \mathbb{E}_{U,W} \left[ \sum_{|\beta|\leq C_{\varepsilon,\delta}} \hat{A}_U(\beta_U)^2 \hat{B}_W(\beta)^2 (1-\varepsilon)^{2|\beta_U|} \right]^{1/2} \quad (4.15)$$

Adding up (4.13), (4.14), and (4.15), we obtain (4.9) (via (4.10)).  $\square$

Now plugging the bounds of Claim 4.9 and Claim 4.10 into (4.4), we see that

$$\frac{\delta^2}{4} \leq \mathbb{E}_{U,W} \left[ \sum_{|\beta|\leq C_{\varepsilon,\delta}} \hat{A}_U(\beta_U)^2 \hat{B}_W(\beta)^2 (1-\varepsilon)^{2|\beta_U|} \right]. \quad (4.16)$$

Note this is similar to the expression we obtained for the Max-3-Lin verifier in Lemma 4.4, with slightly different parameters. Using exactly the same strategy for the two provers with a similar analysis, we obtain a strategy with success probability  $\frac{\delta^4}{64C_{\varepsilon,\delta}}$ .  $\square$

### 4.3 Max-TSA

Thus far we have used Max-TSA with constant, but non-optimal, soundness as the starting point for all of our reductions. Now we also show that even obtaining any non-trivial approximation of Max-TSA is UFG-NP-hard.

**Theorem 4.11.** *For any  $\varepsilon > 0$ , Max-TSA is  $(1, 1/2 + \varepsilon)$ -UFG-NP-hard.*

We now construct the following PCP verifier, which makes queries of the form  $f_{\text{TSA}}(x) = b$ .

**Definition 4.12.** The TSA PCP verifier does the following:

1. Pick a random pair of sets  $(U, W)$  sent to the two provers in the parallel game.
2. Sample  $f \in \mathcal{F}_U$  and  $g_1, g_3, g_4 \in \mathcal{F}_W$  all independently and uniformly at random.
3. Set  $g_2(\mathbf{y}) = f(\mathbf{y}_U) + g_1(\mathbf{y}) + g_3(\mathbf{y}) \wedge g_4(\mathbf{y})$ .
4. Accept if and only if  $A_U^{\text{ff}}(f) + B_W^{\text{ff}}(g_1) + B_W^{\text{ff}}(g_2) + B_W(g_3) \wedge B_W(g_4) = 0$

A subtle yet crucial point is that the queries for  $g_3$  and  $g_4$  are done in the completely unfolded table  $B_W$  (not even folded over true). In this way, the negations introduced by the folding only ever appear on the queries for  $f, g_1$ , and  $g_2$ . Since these queries appear linearly in the test of the verifier, any negations added by the folding can be added up and moved to the right hand side of the test. Thus, the queries of the verifier are indeed all of the form  $f_{\text{TSA}}(\mathbf{x}) = 1$  or  $f_{\text{TSA}}(\mathbf{x}) = 0$  as desired. As usual, [Fact 3.3](#) implies that the reduction given by this PCP verifier is factor graph-preserving. The fact that this PCP has completeness 1 is immediate and again the proof is near-identical to the proof of [Lemma 4.3](#) so we omit it.

**Lemma 4.13.** *The PCP from [Definition 4.12](#) has completeness 1.*

Next we prove soundness.

**Lemma 4.14.** *Let  $\delta > 0$  and set  $C_\delta = \log(\frac{2}{\delta})$ . For  $r > 0$  set  $t = r\frac{4}{\delta^2}C_\delta$ . If the PCP verifier from [Definition 4.12](#) accepts with probability at least  $\frac{1+\delta}{2}$ , then there is a strategy in the original parallel game that makes the verifier accept with probability at least  $\frac{\delta^4}{16C_\delta}$ .*

*Proof.* To emphasize the difference from the folded tables, we write  $C_W = B_W$  throughout this proof for the unfolded table  $B_W$ . If the verifier accepts with probability at least  $\frac{1+\delta}{2}$ , then

$$\begin{aligned} \delta &= \mathbb{E} \left[ (-1)^{A_U^{\text{ff}}(f)} (-1)^{B_W^{\text{ff}}(g_1)} (-1)^{B_W^{\text{ff}}(g_2)} (-1)^{C_W(g_3) \wedge C_W(g_4)} \right] \\ &= \mathbb{E}_{U, W} \left[ \sum_{\alpha, \beta_1, \beta_2} \hat{A}_U^{\text{ff}}(\alpha) \hat{B}_W^{\text{ff}}(\beta_1) \hat{B}_W^{\text{ff}}(\beta_2) \mathbb{E} \left[ \chi_\alpha(f) \chi_{\beta_1}(g_1) \chi_{\beta_2}(g_2) (-1)^{C_W(g_3) \wedge C_W(g_4)} \right] \right] \\ &= \mathbb{E}_{U, W} \left[ \sum_{\alpha, \beta_1, \beta_2} \hat{A}_U^{\text{ff}}(\alpha) \hat{B}_W^{\text{ff}}(\beta_1) \hat{B}_W^{\text{ff}}(\beta_2) \mathbb{E} \left[ \chi_\alpha(f) \chi_{\beta_1}(g_1) \chi_{\beta_2}(f + g_1 + g_3 \wedge g_4) (-1)^{C_W(g_3) \wedge C_W(g_4)} \right] \right] \\ &= \mathbb{E}_{U, W} \left[ \sum_{\beta} \hat{A}_U^{\text{ff}}(\beta_{\oplus U}) \hat{B}_W^{\text{ff}}(\beta)^2 \mathbb{E}_{g_3, g_4} \left[ \chi_\beta(g_3 \wedge g_4) (-1)^{C_W(g_3) \wedge C_W(g_4)} \right] \right]. \end{aligned}$$

Our goal is now to show that the inner expectation over  $g_3$  and  $g_4$  is bounded by a function tending to zero with  $|\beta|$ . First observe that

$$\mathbb{E}_{g_3, g_4} \left[ \chi_\beta(g_3 \wedge g_4) (-1)^{C_W(g_3) \wedge C_W(g_4)} \right]$$

$$= \frac{1}{4} \mathbb{E}_{g_3, g_4} \left[ \chi_\beta(g_3 \wedge g_4) (1 - (-1)^{C_W(g_3)} - (-1)^{C_W(g_4)} + (-1)^{C_W(g_3)} (-1)^{C_W(g_4)}) \right].$$

Observe that by independence of  $g_3$  and  $g_4$ ,

$$\mathbb{E}_{g_3, g_4} [\chi_\beta(g_3 \wedge g_4)] = \mathbb{E}_{g_3} \left[ \prod_{y \in \beta} \mathbb{E}_{g_4} [(-1)^{g_3(y)g_4(y)}] \right] = 2^{-|\beta|}$$

since the product of expectations over  $g_4$  is zero unless  $g_3$  is identically zero on  $\beta$ . Similarly,

$$\left| \mathbb{E}_{g_3, g_4} [\chi_\beta(g_3 \wedge g_4) (-1)^{C_W(g_3)}] \right| = \left| \mathbb{E}_{g_3} \left[ (-1)^{C_W(g_3)} \prod_{y \in \beta} \mathbb{E}_{g_4} [(-1)^{g_3(y)g_4(y)}] \right] \right| \leq 2^{-|\beta|}$$

with the same inequality holding for  $\mathbb{E} [\chi_\beta(g_3 \wedge g_4) (-1)^{C_W(g_4)}]$  by symmetry. Therefore we conclude that

$$\left| \mathbb{E}_{g_3, g_4} [\chi_\beta(g_3 \wedge g_4) (-1)^{C_W(g_3) \wedge C_W(g_4)}] \right| \leq \frac{3}{4} \cdot 2^{-|\beta|} + \frac{1}{4} \left| \mathbb{E}_{g_3, g_4} [\chi_\beta(g_3 \wedge g_4) (-1)^{C_W(g_3)} (-1)^{C_W(g_4)}] \right|.$$

Next we take the Fourier expansion of  $C_W$  in the above expectation and get

$$\mathbb{E}_{g_3, g_4} [\chi_\beta(g_3 \wedge g_4) (-1)^{C_W(g_3)} (-1)^{C_W(g_4)}] = \sum_{\gamma_1, \gamma_2} \hat{C}_W(\gamma_1) \hat{C}_W(\gamma_2) \mathbb{E}_{g_3, g_4} [\chi_\beta(g_3 \wedge g_4) \chi_{\gamma_1}(g_3) \chi_{\gamma_2}(g_4)].$$

Any term above with  $\gamma_1, \gamma_2$  not both being subsets of  $\beta$  has expectation zero. Furthermore, for any choice of  $g_4$  where  $g_4 \wedge \mathbf{1}_\beta \neq \mathbf{1}_{\gamma_1}$ , taking the expectation over  $g_3$  gives us zero, while if  $g_4 \wedge \mathbf{1}_\beta = \mathbf{1}_{\gamma_1}$  the expectation over  $g_3$  equals 1. In the latter case, which happens with probability  $2^{-|\beta|}$  over the choice of  $g_4$ , we have  $\chi_{\gamma_2}(g_4) = \chi_{\gamma_2}(\mathbf{1}_{\gamma_1})$  and thus we are left with

$$2^{-|\beta|} \sum_{\gamma_1, \gamma_2 \subseteq \beta} \hat{C}_W(\gamma_1) \hat{C}_W(\gamma_2) \chi_{\gamma_2}(\mathbf{1}_{\gamma_1}). \quad (4.17)$$

Now observe that for any fixed  $\gamma_1$ ,

$$\left| \sum_{\gamma_2 \subseteq \beta} \hat{C}_W(\gamma_2) \chi_{\gamma_2}(\mathbf{1}_{\gamma_1}) \right| \leq 1,$$

because the sum equals the expected value of  $(-1)^{C_W(g)}$  over all  $g \in \mathcal{F}_W$  which agree with  $\mathbf{1}_{\gamma_1}$  on  $\beta$ . Therefore we can bound (4.17) by

$$2^{-|\beta|} \sum_{\gamma_1 \subseteq \beta} |\hat{C}_W(\gamma_1)| \leq 2^{-|\beta|/2} \sum_{\gamma_1 \subseteq \beta} \hat{C}_W(\gamma_1)^2 \leq 2^{-|\beta|/2}.$$

So we conclude that

$$\left| \mathbb{E}_{g_3, g_4} [\chi_\beta(g_3 \wedge g_4) (-1)^{C_W(g_3) \wedge C_W(g_4)}] \right| \leq \frac{3}{4} \cdot 2^{-|\beta|} + \frac{1}{4} \cdot 2^{-|\beta|/2} < 2^{-|\beta|/2}.$$

Returning to our original formula for the acceptance probability we have by Cauchy-Schwartz and Plancherel

$$\delta^2 \leq \mathbb{E}_{U, W} \left[ \left( \sum_{\beta} |\hat{A}_U^{\boxplus}(\beta_{\oplus U}) \hat{B}_W^{\boxplus}(\beta) 2^{-|\beta|/2}| \right)^2 \right] \leq \mathbb{E}_{U, W} \left[ \sum_{\beta} \hat{A}_U^{\boxplus}(\beta_{\oplus U})^2 \hat{B}_W^{\boxplus}(\beta)^2 2^{-|\beta|} \right]. \quad (4.18)$$

This is essentially the same as the bound (4.1) we obtained in Lemma 4.4 for Max-3-Lin, except with  $2^{-|\beta|}$  in place of  $(1 - 2\varepsilon)^{2|\beta|}$ . We can now use exactly the same strategies for the provers as in that proof and succeed in the smooth parallel repeated game with probability at least  $\frac{\delta^4}{16C_\delta}$ .  $\square$

## 5 Pairwise Independence and Hadamard Predicates

In this section we establish that the results of Chan [Cha16] can be obtained with a universal factor graph. Chan showed that any predicate supporting a pairwise independent subgroup is approximation resistant. In fact, he even showed that such predicates satisfy a strong property called uselessness, introduced by Austrin and Håstad [AH13].

**Definition 5.1.** The predicate  $f : \{0, 1\}^k \rightarrow \{0, 1\}$  is *useless* for a set of functions  $G = \{g : \{0, 1\}^k \rightarrow \mathbb{R}\}$  if for every  $\varepsilon > 0$ , the following promise decision problem is NP-hard. Given a Max-CSP( $f^\pm$ ) instance  $I = (X, C)$ , distinguish between

1. (Yes)  $I$  is  $(1 - \varepsilon)$ -satisfiable.
2. (No) for every assignment and every  $g \in G$  and assignment  $\mathbf{a} : X \rightarrow \{0, 1\}$  it holds that

$$\left| \mathbb{E}_{(S, \mathbf{b}) \sim \mathcal{C}} [g(\mathbf{a}_S + \mathbf{b})] - \mathbb{E}_{\mathbf{u} \sim \{0, 1\}^k} [g(\mathbf{u})] \right| \leq \varepsilon.$$

If  $G$  is the set of all functions on  $k$  bits then we say simply that  $f$  is useless. If the uselessness is established by a factor-graph preserving reduction from any problem that is UFG-NP-hard then we say that  $f$  is *UFG-useless* (for  $G$ ).

The main result of Chan, that we establish with a universal factor graph, can be stated as follows.

**Theorem 5.2** ([Cha16], Theorem 1.1, with a universal factor graph). *Let  $f : \Sigma^k \rightarrow \{0, 1\}$  be any predicate that supports a pairwise independent subgroup of  $\{0, 1\}^k$ . Then  $f$  is UFG-useless.*

As the Hadamard predicates  $\text{Had}_k$  support a pairwise independent subgroup, an immediate and often used corollary is the following.

**Corollary 5.3.** *For any  $\varepsilon > 0$ , Max-Had $_k$  is  $(1 - \varepsilon, (k + 1)2^{-k} + \varepsilon)$ -UFG-NP-hard.*

### 5.1 Analytic Notation, Influences, and Noise

For the purposes of proving Chan's result with universal factor graphs, it turns out to be more notationally convenient to use the  $\{-1, 1\}$  domain for Fourier analysis. Therefore, just for this section, we make the following notational changes. If  $I$  is a set of coordinates  $\mathcal{F}_I$  denotes the set of functions  $f : \{0, 1\}^I \rightarrow \{-1, 1\}$ , and we use  $\mathcal{F}_n$  to denote  $\mathcal{F}_{[n]}$  as before. All of the long code tables  $A$  will be functions  $A : \mathcal{F}_n \rightarrow \{-1, 1\}$ . Finally, for  $\alpha \subseteq \{0, 1\}^n$  the Fourier character  $\chi_\alpha : \mathcal{F}_n \rightarrow \{-1, 1\}$  will be given by

$$\chi_\alpha(f) = \prod_{\mathbf{x} \in \alpha} f(\mathbf{x}).$$

With this notation set up, we additionally need to define influences and the noise operator.

**Definition 5.4.** For a function  $A : \mathcal{F}_n \rightarrow \mathbb{R}$  and set  $\mathcal{B} \subseteq \{0, 1\}^n$  of coordinates, the *influence* of  $\mathcal{B}$  is

$$\text{Inf}_{\mathcal{B}}(A) = \sum_{\alpha \cap \mathcal{B} \neq \emptyset} \hat{A}(\alpha)^2.$$

For a single coordinate  $\mathbf{x} \in \{0, 1\}^n$  we write  $\text{Inf}_{\mathbf{x}}(A)$  for  $\text{Inf}_{\{\mathbf{x}\}}(A)$ .

**Fact 5.5.**  $\text{Inf}_{\mathcal{B}}(A) \leq \sum_{\mathbf{x} \in \mathcal{B}} \text{Inf}_{\mathbf{x}}(A)$

**Definition 5.6.** For a noise rate  $0 \leq \eta \leq 1$  the noise operator  $T_{1-\eta}$  maps functions  $A : \mathcal{F}_n \rightarrow \mathbb{R}$ , to noisy functions  $T_{1-\eta}A : \mathcal{F}_n \rightarrow \mathbb{R}$  defined by

$$T_{1-\eta}A(f) = \mathbb{E}_{\tilde{f} \sim_{1-\eta} f} [A(\tilde{f})],$$

where  $\tilde{f} \sim_{1-\eta} f$  indicates that  $\tilde{f}(\mathbf{x})$  is chosen as  $f(\mathbf{x})$  with probability  $1 - \eta$ , and as a uniformly random bit with probability  $\eta$ , independently for each  $\mathbf{x} \in \{0, 1\}^n$ .

**Fact 5.7.** For every  $\eta > 0$  and every table  $A : \mathcal{F}_n \rightarrow \mathbb{R}$ ,

$$\sum_{\mathbf{x} \in \{0,1\}^n} \text{Inf}_{\mathbf{x}}(T_{1-\eta}A) \leq 1/\eta.$$

## 5.2 Overview

As with other results, we very much follow in the footsteps of the original proof. Given that Chan’s proof is rather long we do not repeat the entire argument here. We only recall some crucial details and describe how to modify them in our setting. The main difference is, not surprisingly, that [Cha16] uses conditioning and this leads both to a simpler proof and the possibility to use simpler notation. To keep notation here simpler, we only present the arguments for the concrete case of Hadamard predicates, but they generalize easily.

The high level view of Chan’s proof is what can be expected. He starts with an instance of label cover with very good soundness. To get better numerical dependencies Chan uses a different starting point, but let us here assume that the starting point is the  $r$ -fold parallel repetition game described in Section 2.4. He then produces a PCP whose acceptance condition is given by  $\text{Had}_k$  and proves that whenever there is a PCP proof where some function  $g$  exceeds its expectation on the answers to a random query, this can be used to derive successful strategies in the two-prover game.

It turns out to be difficult to directly define a PCP where every  $g$  has small expectation. An easier task is to define a PCP where all characters  $\psi : \{-1, 1\}^k \rightarrow \{-1, 1\}$  that are  $j$ -relevant for some fixed  $j \in [k]$  have small expectation. In the Boolean setting, each character is simply a product  $\psi(b_1, \dots, b_k) = \prod_{i \in S} b_i$  for some  $S \subseteq [k]$ , and  $\psi$  is  $j$ -relevant if  $j \in S$ .

We have the following theorem.

**Theorem 5.8** ([Cha16], Theorem 5.4, with a universal factor graph). *For every  $j \in [k]$ ,  $\text{Had}_k$  is UFG-useless for the set of all  $j$ -relevant characters  $\psi : \{-1, 1\}^k \rightarrow \{-1, 1\}$ .*

The proof of this theorem is the main technical part of Chan’s work, and it is also the part that needs modifications in our setting with functional folding in lieu of conditioning. We describe these modifications and the proof in Section 5.3 below.

When we have Theorem 5.8, we can combine it with the very powerful construction, discovered by Chan, of taking the direct sum of instances.

**Definition 5.9.** Given two Max-CSP( $f^\pm$ ) instances  $I = (X, C)$  and  $I' = (X', C')$ , their *direct sum* is defined as  $I \oplus I' = (X \times X', C \oplus C')$ . For each constraint  $f(\mathbf{x}_S + \mathbf{b}) = 1$  in  $C$  and each constraint  $f(\mathbf{x}'_{S'} + \mathbf{b}') = 1$  in  $C'$ , we have the constraint  $f(\mathbf{x}''_{S \oplus S'} + \mathbf{b} + \mathbf{b}') = 1$  in  $C \oplus C'$ , where for two tuples  $\mathbf{u} = (u_1, \dots, u_k)$  and  $\mathbf{v} = (v_1, \dots, v_k)$  we write  $\mathbf{u} \oplus \mathbf{v}$  to denote coordinate-wise concatenation of  $\mathbf{u}$  and  $\mathbf{v}$ , i.e., the tuple of pairs  $((u_1, v_1), \dots, (u_k, v_k))$ .

As shown by Chan ([Cha16], Lemma 5.3), taking direct sum preserves uselessness for characters (if either  $I$  or  $I'$  satisfies the “No” case of Definition 5.1 with respect to some character  $\psi$  then  $I \oplus I'$  does as well). Thus, since the characters form an orthonormal basis for all functions  $g : \{-1, 1\}^k \rightarrow \mathbb{R}$ , taking the direct sum of the  $k$  instances arising from Theorem 5.8 and making the following observation we obtain Theorem 5.2.

*Observation 5.10.* The factor graph of  $I \oplus I'$  depends only on the factor graphs of  $I$  and  $I'$ , and not on the negation patterns in  $I$  and  $I'$ .

### 5.3 Protocol For a Single Coordinate

In this section we sketch Chan’s proof of [Theorem 5.8](#) and the modifications needed to make it hold with a universal factor graph.

Throughout this section, fix the value of the index  $j \in [k]$ , and let  $J = [k] \setminus \{j\}$ , i.e., all elements except  $j$ . Let  $\eta > 0$  be a small parameter to be chosen later.

Given a Max-TSA instance  $I$  we construct a new set of variables  $X'$  as in the reduction template [Definition 3.5](#), and construct the following PCP verifier.

1. Pick a random pair of sets  $(U, W)$  sent to the two provers in the parallel game.
2. Pick a uniformly random function  $f \in \mathcal{F}_U$ .
3. For  $i \in J$  and  $\mathbf{y} \in \{0, 1\}^W$  choose  $g_i(\mathbf{y})$  uniformly at random subject to condition that the string  $(g_i(\mathbf{y}))_{i \in J}$  with  $f(\mathbf{y}_U)$  inserted in position  $j$  satisfies  $\text{Had}_k$ .
4. Let  $\tilde{f}$  and  $\tilde{g}_i$  be  $\eta$ -noisy perturbations of  $f$  and  $g_i$ .
5. Accept if and only if  $(B_W^\boxplus(\tilde{g}_i))_{i \in J}$  with  $A_U^\boxplus(\tilde{f})$  inserted in the  $j$ th position satisfies  $\text{Had}_k$ .

*Remark 5.11.* Let us briefly compare the notation used here to the notation used in Chan’s protocol. The following table shows the notation used for the main objects.

	$j$ 'th coordinate					$i$ 'th coordinate for $i \neq j$				
Chan’s notation	$f_j$	$\tilde{f}_j$	$g_j$	$\mathbf{z}^{(j)}$	$\mathbf{z}_t^{(j)}$	$f_i$	$\tilde{f}_i$	$g_i$	$\mathbf{z}^{(i)}$	$\mathbf{z}_s^{(i)}$
Our notation	$A_U$	$A_U^\boxplus$	$T_{1-\eta}A_U^\boxplus$	$f$	$f(\mathbf{x})$	$B_W$	$B_W^\boxplus$	$T_{1-\eta}B_W^\boxplus$	$g_i$	$g_i(\mathbf{y})$

Note in particular that while the  $g_i$ ’s in Chan’s protocol are the purported long codes with  $\eta$ -noise applied, the  $g_i$ ’s in our protocol are the inputs to the purported long codes on the  $W$  side.

The completeness analysis of the above protocol is easy and is not affected by the modifications we have made, so let us turn to the soundness analysis.

Fix some  $j$ -relevant character  $\psi$  and suppose that the expectation of  $\psi$  over the answers to the provers deviates from its expectation (0) by at least  $\varepsilon$ , i.e.,

$$\left| \mathbb{E}_{(U,W)} \left[ \mathbb{E}_{\tilde{f}, \{\tilde{g}_i\}_{i \in J}} [\psi(A_U^\boxplus(\tilde{f}), \{B_W^\boxplus(\tilde{g}_i)\}_{i \in J})] \right] \right| > \varepsilon.$$

By Markov’s inequality, it holds that for at least an  $\varepsilon/2$  fraction of all query pairs  $(U, W)$ , the inner expectation is at least  $\varepsilon/2$  in absolute value. Fix one such “good” pair  $(U, W)$  and to simplify notation let  $A(f) = T_{1-\eta}A_U^\boxplus$  and  $B(g) = T_{1-\eta}B_W^\boxplus$ . Thus we have

$$\left| \mathbb{E}_{f, \{g_i\}_{i \in J}} [\psi(A(f), \{B(g_i)\}_{i \in J})] \right| > \varepsilon/2. \quad (5.1)$$

For  $\mathbf{x} \in \{0, 1\}^U$ , define  $\mathcal{B}(\mathbf{x})$  to be the set of  $\mathbf{y} \in \{0, 1\}^W$  such that  $\mathbf{y}_U = \mathbf{x}$  (in Chan’s notation, this is the “block”  $B(t)$ ). The key quantity to study is

$$\sum_{\mathbf{x}} \text{Inf}_{\mathbf{x}}(A) \text{Inf}_{\mathcal{B}(\mathbf{x})}(B), \quad (5.2)$$

which measures the presence of common influences between the noised tables  $A$  and  $B$ . Using an invariance-style proof, it is shown that when (5.1) holds then (5.2) must also be large. Concretely we have the following theorem.

**Theorem 5.12** ([Cha16], Theorem 6.7). *In the notation above, let  $\mathcal{Z} \subseteq \{0,1\}^U$  be any set of assignments such that*

$$\sum_{\mathbf{x} \notin \mathcal{Z}} \text{Inf}_{\mathbf{x}}(A) \text{Inf}_{\mathcal{B}(\mathbf{x})}(B) \leq \tau,$$

and define  $A^{\mathcal{Z}} : \mathcal{F}_U \rightarrow [0,1]$  to be the part of  $A$  depending only on  $\mathcal{Z}$ .<sup>2</sup> Then

$$\left| \mathbb{E}_{f, \{g_i\}_{i \in J}} [\psi(A(f), \{B(g_i)\}_{i \in J})] \right| \leq \left| \mathbb{E}_{f, \{g_i\}_{i \in J}} [\psi(A^{\mathcal{Z}}(f), \{B(g_i)\}_{i \in J})] \right| + \delta(k, \eta, \tau)$$

where for every fixed  $k$  and  $\eta$ ,  $\delta(k, \eta, \tau)$  tends to 0 as  $\tau$  tends to 0.

This is a theorem purely about analysis of Boolean functions and its proof relies only on the pairwise independence of the underlying CSP and not on the structure of the parallel repeated game, and as such it applies without modification in our setting. Chan only states the theorem for  $\mathcal{Z} = \emptyset$  but inspection of the proof, which is based on rerandomizing one coordinate at a time, reveals that it holds for any  $\mathcal{Z}$ .

In Chan's original proof, the case  $\mathcal{Z} = \emptyset$  is all that is needed, since combined with (5.1) it lets us conclude that the tables  $A$  and  $B$  have shared influential coordinates, which can then immediately be used in a standard way to define strategies that are accepted with constant probability in the parallel repeated game.

In our setting it is not a priori clear that this yields a good strategy, since the functional folding might not guarantee that any influential coordinate of  $A$  or  $B$  actually satisfies the constraints. However, as we shall now see, it turns out that this is indeed the case, so the same strategy does work also in our setting<sup>3</sup>.

Since  $A$  and  $B$  are tables with noise applied, it follows from Fact 5.7 and Fact 5.5 that if we let  $\mathcal{Z}$  be the set of assignments  $\mathbf{x}$  such that both

$$\text{Inf}_{\mathbf{x}}(A) \geq \tau\eta/2 \quad \text{and} \quad \text{Inf}_{\mathcal{B}(\mathbf{x})}(B) \geq \tau\eta/2 \quad (5.3)$$

then

$$\sum_{\mathbf{x} \notin \mathcal{Z}} \text{Inf}_{\mathbf{x}}(A) \text{Inf}_{\mathcal{B}(\mathbf{x})}(B) \leq \tau.$$

Choosing  $\tau$  small enough so that  $\delta(k, \eta, \tau) < \varepsilon/2$ , it follows from Theorem 5.12 and (5.1) that

$$\left| \mathbb{E}_{f, \{g_i\}_{i \in J}} [\psi(A^{\mathcal{Z}}(f), \{B(g_i)\}_{i \in J})] \right| > 0 \quad (5.4)$$

We now use the functional folding, and have the following observation.

*Claim 5.13.* If no  $\mathbf{x} \in \mathcal{Z}$  satisfies all equations  $h_i(\mathbf{x}) = b_i$  for  $i = 1, \dots, t$ , then  $A^{\mathcal{Z}}$  is identically 0.

*Proof.* By Lemma 3.4 each non-zero Fourier coefficient of  $A$  contains an element that satisfies all the constraints. If we rerandomize all these values then the expectation is zero.  $\square$

Since  $\psi$  is a  $j$ -relevant character, the left hand side of (5.4) would be 0 if  $A^{\mathcal{Z}}$  was identically 0, so it follows that there must be some  $\mathbf{x}^* \in \mathcal{Z}$  which satisfies all equations on  $U$ .

We now define the strategies for the provers in the repeated game in a standard way: we choose answer  $\mathbf{x}$  for  $U$  with probability proportional to  $\text{Inf}_{\mathbf{x}}(A)$  and, independently, analogously for  $W$ . By Fact 5.7,  $\mathbf{x}$  and  $\mathbf{y}$  are chosen with probabilities at least  $\eta \text{Inf}_{\mathbf{x}}(A)$  and  $\eta \text{Inf}_{\mathbf{y}}(B)$ , so the probability

<sup>2</sup>Equivalently,  $A^{\mathcal{Z}}(f)$  is the expectation of  $A(\bar{f})$  on a copy of  $f$  where all coordinates outside  $\mathcal{Z}$  have been rerandomized.

<sup>3</sup>Of course, there is no reason why the provers would ever output a value that does not satisfy the relevant constraints but it is slightly easier to analyze this variant of their strategies.

that the good assignment  $\mathbf{x}^*$  is chosen and is consistent with the answer of the other prover is at least

$$\eta \operatorname{Inf}_{\mathbf{x}^*}(A) \sum_{\mathbf{y} \in \mathcal{B}(\mathbf{x}^*)} \eta \operatorname{Inf}_{\mathbf{y}}(B) \geq \eta^2 \operatorname{Inf}_{\mathbf{x}^*}(A) \operatorname{Inf}_{\mathcal{B}(\mathbf{x}^*)}(B) \geq \eta^4 \tau^2 / 4.$$

Aggregating this over the  $\varepsilon/2$  fraction of good pairs  $(U, W)$  of queries in the repeated game, we conclude that in expectation a fraction  $\varepsilon \eta^4 \tau^2 / 8$  of all query pairs are assigned answers that are consistent and where all constraints on  $U$  are satisfied. The only remaining issue is to establish that the answers of the other prover often satisfy the additional constraints on  $W$ .

Fix any assignment to the variables in  $W$  which does not satisfy all constraints. By [Claim 2.20](#) the probability, over the choice of  $U$ , that it projects to an assignment that satisfies all constraints on this smaller set is bounded by  $r/(r+t)$ . It follows that if we choose  $t \geq 16r\varepsilon^{-1}\eta^{-4}\tau^{-2}$  then the total expected fraction of query pairs where the assignments are consistent and satisfy the constraints on  $U$  but not those on  $W$  is bounded by  $\frac{r}{r+t} \leq \varepsilon \eta^4 \tau^2 / 16$ . We conclude that in expectation the influence-based random strategy wins the repeated game with probability at least  $\varepsilon \eta^4 \tau^2 / 16$ . This concludes the description of the modifications of the proof Chan and our proof of [Theorem 5.8](#).

## 6 Promise CSPs

Functional folding can also be used to obtain hardness results for *promise CSPs* (PCSPs) with universal factor graphs. We first recall the pertinent definitions.

**Definition 6.1.** A *PCSP language* is a pair  $(\Gamma, \Lambda)$  of two indexed constraint languages  $\Gamma = \{f_1, \dots, f_t\}$  and  $\Lambda = \{g_1, \dots, g_t\}$  such that  $f_i$  and  $g_i$  have the same arity and  $f_i(\mathbf{x}) \leq g_i(\mathbf{x})$  for all  $i$  and  $x$ .

A PCSP language has *free negations*, if for every  $k$ -ary constraint pair  $(f_i, g_i)$  and every  $\mathbf{b} \in \{0, 1\}^k$ , the constraint pair  $(f_i^{\mathbf{b}}, g_i^{\mathbf{b}})$  is also in the language, where  $f_i^{\mathbf{b}}(\mathbf{x}) = f_i(\mathbf{x} + \mathbf{b})$ .

An instance  $I$  of the PCSP $(\Gamma, \Lambda)$  problem is a pair  $(X, C)$  where  $X$  is a set of variables and  $C$  a set of constraints. Each constraint  $c \in C$  is a pair  $(i, S)$ , for a constraint type  $i \in [t]$  and scope  $S$ . We write  $I_\Gamma$  for the induced CSP $(\Gamma)$  instance where each constraint  $(i, S)$  is replaced by  $(f_i, S)$  and  $I_\Lambda$  for the induced CSP $(\Lambda)$  instance where  $(i, S)$  is replaced by  $(g_i, S)$ .

PCSP $(\Gamma, \Lambda)$  is the promise decision problem where given an instance  $I$  the objective is to distinguish whether  $I_\Gamma$  is satisfiable or  $I_\Lambda$  is unsatisfiable.

A prototypical example of a PCSP (with free negations) is the “ $(2 + \varepsilon)$ -Sat” problem, in which we are given a  $(2k + 1)$ -CNF formula  $\phi$  and the objective is to distinguish the case where there is an assignment satisfying at least  $k$  literals in every clause of  $\phi$ , from the case where  $\phi$  is unsatisfiable.

[Definition 2.12](#) of UFG-NP-hardness extends naturally to PCSP problems. To state the hardness result for PCSPs, we also need the notion of polymorphisms, defined next.

**Definition 6.2.** A *polymorphism* of a PCSP language  $(\Gamma, \Lambda)$  is a function  $p : \Sigma^\ell \rightarrow \Sigma$  such that, for every pair of constraint types  $(f_i, g_i) \in (\Gamma, \Lambda)$  and all  $x_1, \dots, x_n \in f_i^{-1}(1)$  (where  $k$  is the arity of  $f_i$  and  $g_i$ ) it holds that

$$(p(x_{1,1}, \dots, x_{\ell,1}), p(x_{1,2}, \dots, x_{\ell,2}), \dots, p(x_{1,k}, \dots, x_{\ell,k})) \in g_i^{-1}(1).$$

A polymorphism  $p$  is *folded* if  $p(\neg \mathbf{x}) = \neg p(\mathbf{x})$  for all  $\mathbf{x} \in \{0, 1\}^\ell$ .

Our main hardness result of PCSPs having universal factor graphs is the following.

**Theorem 6.3.** *Let  $(\Gamma, \Lambda)$  be a finite PCSP language with free negations, and suppose that there exists a universal constant  $C = C(\Gamma, \Lambda) < \infty$  such that every folded polymorphism of  $(\Gamma, \Lambda)$  is a  $C$ -junta. Then PCSP $(\Gamma, \Lambda)$  is UFG-NP-hard.*

This theorem is approximately Theorem 4.7 of [AGH17]. That theorem was simplified and generalized by Brakensiek and Guruswami [BG18] to PCSP languages where the polymorphisms are only required to be  $C$ -fixing, a weaker condition than being a  $C$ -junta where it is only required that setting all the  $C$  coordinates to 0 fixes the value of the function. Another difference is that the result of Brakensiek and Guruswami does not require the PCSP language to have free negations, which is something we require in order to be able to apply functional folding. The proof below follows the simplified proof of Brakensiek and Guruswami very closely but there is one step where we need the stronger condition of being a  $C$ -junta instead of just  $C$ -fixing—see [Footnote 4](#) for further details.

Using the fact that the polymorphisms of  $(2 + \varepsilon)$ -Sat are juntas [AGH17], we have the following immediate corollary.

**Corollary 6.4.**  $(2 + \varepsilon)$ -Sat is UFG-NP-hard.

*Proof of Theorem 6.3.* Given a Max-TSA instance  $I$  we construct a new set of variables  $X'$  as in the reduction template [Definition 3.5](#).

For every query  $U$  (resp.  $W$  to  $P_2$ ) in the repeated game, we add constraints on  $A_U^{\boxplus}$  (resp.  $B_W^{\boxplus}$ ) forcing it to be a polymorphism of  $(\Gamma, \Lambda)$ .

Furthermore, for every pair of queries  $U \subseteq W$  sent to the two provers in the parallel game, and all functions  $f \in \mathcal{F}_U$ , we identify the values of  $A_U^{\boxplus}(f)$  and  $B_W^{\boxplus}(f)$  (where we think of  $f \in \mathcal{F}_U$  as a function  $f \in \mathcal{F}_W$  that only depends on the coordinates in  $U$ , in the obvious way). This simply means that whenever we would have accessed  $A_U^{\boxplus}(f)$ , we instead access  $B_W^{\boxplus}(f)$ . It is clear that this construction is factor graph-preserving.

*Claim 6.5 (Completeness).* If  $I$  is satisfiable then  $R(I)_\Gamma$  is satisfiable.

The proof of completeness is immediate from the definitions and we omit it.

*Claim 6.6 (Soundness).* If  $R(I)_\Lambda$  is satisfiable and  $t \geq C^2 r$  then  $I$  is  $\frac{1}{2C}$ -satisfiable.

*Proof.* Given a satisfying assignment (consisting of supposed long codes) to  $R(I)_\Lambda$ , let  $\alpha^U$  (resp.  $\beta^W$ ) be the set of up to  $C$  coordinates that  $A_U^{\boxplus}$  (resp.  $B_W^{\boxplus}$ ) depends on.

The key observation, that we now proceed to establish, is that for a pair of queries  $U \subseteq W$  sent to the two provers such that  $|\beta_U^W| = |\alpha^U|$  (which, by the choice  $t \geq C^2 r$  and [Claim 2.20](#), are at least  $1/2$  of all query pairs), we must have  $\beta_U^W \subseteq \alpha^U$ . Indeed, suppose for contradiction that  $\mathbf{x}^* \in \beta^W$  but  $\mathbf{x}_U^* \notin \alpha^U$ . Let  $g \in \mathcal{F}_W$  be a function such that  $B_W^{\boxplus}(g) \neq B_W^{\boxplus}(g + \mathbf{1}_{\mathbf{x}^*})$  and  $g(\mathbf{x}) = 1$  for all  $\mathbf{x} \notin \beta^W$ . Because  $|\beta_U^W| = |\alpha^U|$  we can also view  $g$  as a function  $f \in \mathcal{F}_U$  (defined by  $f(\mathbf{x}) = 1$  if  $\mathbf{y} \notin \beta_U^W$  and otherwise  $f(\mathbf{x}) = g(\mathbf{y})$  where  $\mathbf{y}$  is the unique  $\mathbf{y} \in \beta^W$  such that  $\mathbf{y}_U = \mathbf{x}$ ). Then using that  $\mathbf{x}_U^* \notin \alpha^U$  and the identification of values in  $A_U$  and  $B_W$  we have the contradiction

$$B_W(g + \mathbf{1}_{\mathbf{x}^*}) = A_U(f + \mathbf{1}_{\mathbf{x}_U^*}) = A_U(f) = B_W(f) \neq B_W(f + \mathbf{1}_{\mathbf{x}^*}) = B_W(g + \mathbf{1}_{\mathbf{x}^*})$$

and the key observation follows.<sup>4</sup>

By [Lemma 3.4](#), at least one  $x \in \beta^W$  satisfies all constraints in  $W$ , and the strategy for  $P_2$  in the repeated game is to use an arbitrary such  $x$ . The strategy for  $P_1$  is to select a random assignment  $x \in \alpha^U$ . Since at least half the query pairs satisfy  $|\beta_U^W| = |\alpha^U|$  and  $|\alpha^U| \leq C$  for all  $U$ , this strategy is accepted with probability at least  $\frac{1}{2C}$ .  $\square$

Combining the completeness and soundness claims, the theorem follows.  $\square$

---

<sup>4</sup> This argument is where we use that the polymorphisms are  $C$ -juntas as opposed to just  $C$ -fixing. In the previous PCSP hardness results it was sufficient to establish that  $\beta_U^W \cap \alpha^U \neq \emptyset$ , but this is not sufficient for us. In our setting the functional folding only guarantees that at least one  $\mathbf{x} \in \beta^W$  satisfies the equations on  $W$  (and this property is easy to establish also for tables that are  $C$ -fixing rather than  $C$ -juntas), and we need to make sure that this specific  $\mathbf{x}$  projects to an element of  $\alpha^U$ . For this reason, just having non-empty intersection between  $\beta_U^W$  and  $\alpha^U$  is not sufficient.

## 7 Miscellaneous Extensions

In this section we discuss various further extensions to our results.

### 7.1 More Hardness Results by Gadgets

One major method for deriving new hardness results is by a method usually referred to as “gadget reductions”. In such a reduction from  $\text{Max-CSP}(\Gamma_1)$  to  $\text{Max-CSP}(\Gamma_2)$  one takes one constraint in the source problem and produces one or several constraints in the target problem. These new constraints contain the variables from the original problem and some new variables which are unique to the constraint processed. A general theory for constructing optimal gadgets was introduced by Trevisan et al. [TSSW00].

To make such a reduction factor graph-preserving one simply needs to ensure that the factor graph of constant size obtained from a single constraint does not depend on which constraint from the family  $\Gamma_1$  was used. This is a simple property to test and turns out to be true for most reductions. In particular, one favorite starting point of such a reduction is Max-3-Lin and if we allow negations of variables then  $\Gamma_1$  is the single predicate parity. This implies that as soon as we reduce to another Max-CSP that allows negations, the reduction is automatically factor graph-preserving. Let us state a couple of immediate corollaries to this observation and some reductions described in [Hås01] and constructed based on the methods of [TSSW00].

**Corollary 7.1.** *For any  $\varepsilon > 0$ , Max-2-Lin is  $(\frac{3}{4} - \varepsilon, \frac{11}{16} + \varepsilon)$ -UFG-NP-hard.*

The reduction takes a single equation of the form  $x + y + z = 0$  and produces 16 equations each containing two variables from the set  $\{x, y, z\}$  joint with 5 new variables. If the equation is satisfied then we can set the new variables to satisfy 12 equations while if it is not satisfied it is only possible to satisfy 10 equations. For 2-Sat we have the following corollary.

**Corollary 7.2.** *For any  $\varepsilon > 0$  Max-2-Sat is  $(\frac{11}{12} - \varepsilon, \frac{21}{24} + \varepsilon)$ -UFG-NP-hard.*

Here the reduction takes one equation and produces 12 clauses of size two of which 11 can be satisfied if the equation is satisfied while if it is not, the optimum is 10.

It might be instructive to see what happens to the gadget reduction from Max-3-Lin to Max-Cut in [TSSW00]. Here each variable in the original problem corresponds to a node in the resulting Max-Cut instance. Each equation containing  $x, y$  and  $z$  produces a set of new variables which are connected in one of two ways depending on whether the right hand side of the equation is 0 or 1. This implies that the reduction is not factor graph-preserving. Of course, this is not very surprising since when Max-Cut is viewed as a Max-CSP the predicate family is the single predicate of non-equality. This implies that the factor graph uniquely defines the instance and thus it is not an interesting problem in the current context.

### 7.2 Larger Domains

Several of the hardness results that we reproved with universal factor graphs are known to apply also to CSPs with larger domains (e.g., the Max-3-Lin mod  $q$  problem). While we have chosen to focus on Boolean CSPs throughout the paper to keep the notation as simple as possible and focus on the core ideas, these results for larger domains can also be obtained in the universal factor graph setting.

In particular, the approximation resistance of Max-3-Lin (Theorem 4.1) and Max-TSA (Theorem 4.11), and the uselessness of predicates supporting pairwise independent subgroups (Theorem 5.2) can be generalized to arbitrary domains. Let us briefly describe how. Here we do not go into depth and assume to a greater extent than in other parts of the paper that the reader is familiar with the corresponding results in the standard settings, and the Fourier transform over  $\mathbb{Z}_q$ .

For a domain of size  $q$ , the long code-based reductions from the parallel repeated game are modified in the exact same way as one modifies the standard hardness reductions for these problems,

by working with  $q$ -ary long code tables  $A_U : \mathbb{Z}_q^U \rightarrow \mathbb{Z}_q$  and  $B_W : \mathbb{Z}_q^W \rightarrow \mathbb{Z}_q$  and doing Fourier analysis over  $\mathbb{Z}_q^n$  instead of over  $\{0, 1\}^n$ . However, we can not start from a parallel repetition of the Max-TSA problem, but instead need to start over some system of equations over  $\mathbb{Z}_q$ . This is because the functional folding involves taking quotients of the domain  $\mathbb{Z}_q^U$  over the constraint equations of our CSP instance and hence those constraint equations also need to be over  $\mathbb{Z}_q$  rather than  $\mathbb{Z}_2$ .

Fortunately, there are several easy ways to overcome this obstacle to obtain a starting point that can be used and we now sketch one. Starting with a Max-3-Sat instance  $I$ , we construct the following system of equations over  $\mathbb{Z}_q$ . For each clause  $x^a \vee y^b \vee z^c$ , where  $x, y$ , and  $z$  are variables, and  $a, b, c \in \{-1, 1\}$  indicate whether a variable appears positively or negatively, add three new variables  $X, Y, Z$  (separately for each clause of  $I$ ) and four equations

$$o(X, Y, Z) = 1 \quad x \cdot X = a \quad y \cdot Y = b \quad z \cdot Z = c$$

over  $\mathbb{Z}_q$ , where the function  $o(X, Y, Z)$  is 1 if and only if at least one of  $X, Y$ , and  $Z$  equals  $-1$ . It is easy to see that if the best assignment to  $I$  falsifies a  $\delta$  fraction of clauses then the best assignment to the system of equations falsifies a  $\delta/4$  fraction of equations.

Furthermore, the left hand sides of the equations depend only on the factor graph of  $I$ , and the negations of  $I$  only appear as right hand sides of equations. Thus this is a factor graph-preserving reduction establishing  $(1, 1 - \delta)$ -UFG-NP-hardness for an equational<sup>5</sup> Max-CSP over  $\mathbb{Z}_q$ . So analogously to the Boolean case, we apply smooth parallel repetition, introduce  $q$ -ary long codes and apply functional folding to these. We then have the following analogue over [Lemma 3.4](#) which says that any non-zero Fourier coefficient of a functionally folded table must depend on an assignment satisfying all the constraints.

**Lemma 7.3.** *Let  $A$  be a supposed  $q$ -ary long code,  $\{h_i(\mathbf{x}) = b_i\}_{i=1}^r$  be a set of equational constraints over  $\mathbb{Z}_q$ ,  $\mathbf{b} = (b_1, \dots, b_r) \in \{0, 1\}^r$ , and  $H = \{h_1, \dots, h_r\}$ . Let  $A_{H,b}$  be  $A$  folded over all functions on  $H$  with respect to  $\mathbf{b}$ . If  $\hat{A}_{H,b}(\alpha) \neq 0$  then the sum of  $\alpha(\mathbf{x})$  over the assignments  $\mathbf{x}$  that satisfy all  $r$  equations equals  $1 \pmod q$ .*

*Proof.* Let  $h(\mathbf{x}) = (h_1(\mathbf{x}), \dots, h_r(\mathbf{x}))$ . Recall that

$$\hat{A}_{H,b}(\alpha) = \mathbb{E}_f \left[ \omega^{A_{H,b}(f) - \langle \alpha, f \rangle} \right],$$

where  $\omega = e^{2\pi i/q}$  is a complex  $q$ 'th root of unity and  $\langle \alpha, f \rangle = \sum_{\mathbf{x}} \alpha(\mathbf{x})f(\mathbf{x})$  is the inner product of the functions  $\alpha$  and  $f$ .

By the folding we have for every  $c \in \mathbb{Z}_q$  that

$$A_{H,b}(f + c \cdot \mathbf{1}_b(h)) = A_{H,b}(f) + c \cdot \mathbf{1}_b(\mathbf{b}) = A_{H,b}(f) + c.$$

Let  $z$  be the sum of  $\alpha(\mathbf{x})$  over all  $\mathbf{x}$  satisfying all  $r$  equations. Then

$$\omega^{\langle \alpha, f + c\mathbf{1}_b(h) \rangle} = \omega^{\langle \alpha, f \rangle + c \sum_{\mathbf{1}_b(h(\mathbf{x}))=1} \alpha(\mathbf{x})} = \omega^{\langle \alpha, f \rangle + cz}.$$

Since the distribution over  $f + c\mathbf{1}_b(h)$  over randomly chosen  $f$  and  $c$  is the same as the distribution over  $f$ , it follows that

$$\hat{A}_{H,b}(\alpha) = \mathbb{E}_{f,c} \left[ \omega^{A_{H,b}(f) + c - \langle \alpha, f \rangle - cz} \right] = \hat{A}_{H,b}(\alpha) \cdot \mathbb{E}_c \left[ \omega^{c(1-z)} \right].$$

The expectation over  $c$  is 0 unless  $z = 1$ , so the claim follows.  $\square$

With this key property of functional folding established, it is a straightforward but tedious task to go over the existing hardness of approximation proofs for these results and adapt them to the universal factor graph setting in exactly the same way as done for the Boolean case in the preceding sections.

<sup>5</sup>For an appropriate generalization of equational CSPs to the setting where we have two different types of equations instead of just one as in [Definition 2.4](#).

## 8 Concluding Remarks and Open Questions

We have established that many of the current best inapproximability results for various Max-CSPs and PCSPs can be made to hold with universal factor graphs, meaning that the hardness of the problems stem from the variable negations and not from the constraint-variable incidence structure.

Given these new hardness results one can wonder whether there are any natural situations where preprocessing helps. As discussed in the introduction, Max-3-Lin in the universal factor graph setting corresponds to having a fixed linear code and the input is only the vector to which one wants to find a close point. In a similar situation, where one is given a fixed integer lattice and asked to find close points to input vectors, preprocessing seems to help [LLS90, DRS14], but we do not know of a corresponding result for the problem on codes.

The only natural example we are aware of where preprocessing seems to help in the CSP setting is the example [FKO06] pointed out in [FJ12]. Here a graph structure in the the factor graph can be used to efficiently refute random instances of 3-Sat with  $n^{1.4}$  clauses. In addition to this one can come up with contrived examples, for instance by taking a Max-CSP consisting of only two predicates, one being very sparse and hard to approximate and one being very dense. Then without universal factor graphs this problem is very hard to approximate due to the sparse predicate, but in the universal factor graph setting an algorithm can precompute the optimal solution to the instance when all constraints use the sparse predicate, and then either use this assignment or a random assignment to get a better approximation ratio.

The landscape of CSPs would certainly be more interesting if preprocessing was helpful in more general situations, and as we mention below there are some natural problems where we currently do not know whether preprocessing helps or not.

Let us mention some interesting avenues for potential future work.

1. From an efficiency point of view our reductions leave something to be desired. The main source of this is our need to use smooth parallel repetition, which incurs a large polynomial blow-up with the degree depending on  $\varepsilon$ . E.g., as a consequence our results do not rule out approximating Max-3-Lin with factor graph preprocessing within a factor  $1/2 + \varepsilon$  in time  $\exp(n^\varepsilon)$ , whereas approximating Max-3-Lin *without* factor graph preprocessing to within  $1/2 + o(1)$  does not even have  $\exp(n^{1-o(1)})$  time algorithms assuming ETH [MR08].
2. All our hardness results are based on functional folding, which inherently introduces negations, either through negated literals or through linearity as in Max-3-Lin and Max-TSA. As such, our methods can not be used to prove hardness of problems where folding over true is not possible, such as the Max-3-Sat problem without mixed clauses, which is known to be NP-hard to approximate within  $7/8 + \varepsilon$  (even on satisfiable instances) [GK05]. Does this and similar problems remain equally hard to approximate with factor graph preprocessing?
3. It would be interesting to obtain universal factor graphs for problems whose hardness is based on the Unique Games Conjecture (UGC), such as approximating Max-2-Lin to within a factor 0.879 [KKMO07] or Max-2-Sat to within a factor 0.941 [Aus07]. Any such result would probably have to be based on some strengthened version of the UGC, but it is very unclear to us even what a suitable formulation of such a strengthened UGC could be that would allow us to perform the reduction to Max-2-Lin in a factor graph-preserving way.

## References

- [AGH17] Per Austrin, Venkatesan Guruswami, and Johan Håstad, *(2+ $\varepsilon$ )-sat is np-hard*, SIAM J. Comput. **46** (2017), no. 5, 1554–1573.
- [AH13] Per Austrin and Johan Håstad, *On the usefulness of predicates*, TOCT **5** (2013), no. 1, 1:1–1:24.

- [Aus07] Per Austrin, *Balanced max 2-sat might not be the hardest*, Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007, 2007, pp. 189–197.
- [BG18] Joshua Brakensiek and Venkatesan Guruswami, *Promise constraint satisfaction: Structure theory and a symmetric boolean dichotomy*, Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018, 2018, pp. 1782–1801.
- [BGS98] M. Bellare, O. Goldreich, and M. Sudan, *Free bits, PCPs and non-approximability—towards tight results.*, SIAM Journal on Computing **27** (1998), 804–915.
- [BKO19] Jakub Bulín, Andrei A. Krokhin, and Jakub Oprsal, *Algebraic approach to promise constraint satisfaction*, Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019., 2019, pp. 602–613.
- [Cha16] Siu On Chan, *Approximation resistance from pairwise independent subgroups*, Journal of the ACM **63** (2016), 1–32.
- [DKV02] Víctor Dalmau, Phokion G. Kolaitis, and Moshe Y. Vardi, *Constraint satisfaction, bounded treewidth, and finite-variable logics*, Principles and Practice of Constraint Programming - CP 2002 (Berlin, Heidelberg) (Pascal Van Hentenryck, ed.), Springer Berlin Heidelberg, 2002, pp. 310–326.
- [DRS14] Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz, *On the closest vector problem with a distance guarantee*, IEEE 29th Conference on Computational Complexity, CCC 2014, Vancouver, BC, Canada, June 11-13, 2014, 2014, pp. 98–109.
- [FJ12] Uriel Feige and Shlomo Jozeph, *Universal factor graphs*, Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I (Artur Czumaj, Kurt Mehlhorn, Andrew M. Pitts, and Roger Wattenhofer, eds.), Lecture Notes in Computer Science, vol. 7391, Springer, 2012, pp. 339–350.
- [FKO06] Uriel Feige, Jeong Han Kim, and Eran Ofek, *Witnesses for non-satisfiability of dense random 3cnf formulas*, Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (Washington, DC, USA), FOCS '06, IEEE Computer Society, 2006, pp. 497–508.
- [GK05] Venkatesan Guruswami and Subhash Khot, *Hardness of max 3sat with no mixed clauses*, 20th Annual IEEE Conference on Computational Complexity (CCC 2005), 11-15 June 2005, San Jose, CA, USA, 2005, pp. 154–162.
- [Gol11] Oded Goldreich, *Candidate one-way functions based on expander graphs*, pp. 76–87, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [Gro07] Martin Grohe, *The complexity of homomorphism and constraint satisfaction problems seen from the other side*, J. ACM **54** (2007), no. 1, 1:1–1:24.
- [Hås01] Johan Håstad, *Some optimal inapproximability results*, Journal of the ACM **48** (2001), no. 4, 798–859.
- [HKLT19] Prahladh Harsha, Subhash Khot, Euiwoong Lee, and Devanathan Thiruvengatachari, *Improved 3LIN Hardness via Linear Label Cover*, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019) (Dagstuhl, Germany) (Dimitris Achlioptas and László A. Végh, eds.), vol. 145, 2019, pp. 9:1–9:16.

- [Joz14] Shlomo Jozeph, *Universal Factor Graphs for Every NP-Hard Boolean CSP*, Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2014) (Dagstuhl, Germany) (Klaus Jansen, José D. P. Rolim, Nikhil R. Devanur, and Cristopher Moore, eds.), Leibniz International Proceedings in Informatics (LIPIcs), vol. 28, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2014, pp. 274–283.
- [Kho02] Subhash Khot, *Hardness results for coloring 3-colorable 3-uniform hypergraphs*, 43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings, 2002, pp. 23–32.
- [KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell, *Optimal inapproximability results for MAX-CUT and other 2-variable csps?*, SIAM J. Comput. **37** (2007), no. 1, 319–357.
- [LLS90] J. C. Lagarias, H. W. Lenstra, and C. P. Schnorr, *Korkin-zolotarev bases and successive minima of a lattice and its reciprocal lattice*, Combinatorica **10** (1990), no. 4, 333–348.
- [MR08] Dana Moshkovitz and Ran Raz, *Two-query pcp with subconstant error*, J. ACM **57** (2008), no. 5, 29:1–29:29.
- [Raz98] R. Raz, *A parallel repetition theorem*, SIAM J. on Computing **27** (1998), 763–803.
- [TSSW00] L. Trevisan, G. Sorkin, M. Sudan, and D. Williamson, *Gadgets, approximation and linear programming.*, SIAM Journal on Computing **29** (2000), 2074–2097.