

Limitations of Local Quantum Algorithms on Random Max- k -XOR and Beyond

Chi-Ning Chou* Peter J. Love† Juspreet Singh Sandhu‡ Jonathan Shi§

February 23, 2022

Abstract

We introduce a notion of *generic local algorithm* which strictly generalizes existing frameworks of local algorithms such as *factors of i.i.d.* by capturing local *quantum* algorithms such as the Quantum Approximate Optimization Algorithm (QAOA).

Motivated by a question of Farhi et al. [arXiv:1910.08187, 2019] we then show limitations of generic local algorithms including QAOA on random instances of constraint satisfaction problems (CSPs). Specifically, we show that any generic local algorithm whose assignment to a vertex depends only on a local neighborhood with $o(n)$ other vertices (such as the QAOA at depth less than $\epsilon \log(n)$) cannot arbitrarily-well approximate boolean CSPs if the problem satisfies a geometric property from statistical physics called the coupled overlap-gap property (OGP) [Chen et al., Annals of Probability, 47(3), 2019]. We show that the random MAX- k -XOR problem has this property when $k \geq 4$ is even by extending the corresponding result for diluted k -spin glasses.

Our concentration lemmas confirm a conjecture of Brandao et al. [arXiv:1812.04170, 2018] asserting that the landscape independence of QAOA extends to logarithmic depth—in other words, for every fixed choice of QAOA angle parameters, the algorithm at logarithmic depth performs almost equally well on almost all instances.

One of these concentration lemmas is a strengthening of McDiarmid’s inequality, applicable when the random variables have a highly biased distribution, and may be of independent interest.

*School of Engineering & Applied Sciences, Harvard University, Cambridge, Massachusetts, USA. Supported by NSF awards CCF 1565264 and CNS 1618026. Email: chiningchou@g.harvard.edu.

†Department of Physics & Astronomy, Tufts University, Medford, Massachusetts, USA. Supported by DARPA ONISQ program award HR001120C0068. Email: peter.love@tufts.edu.

‡School of Engineering & Applied Sciences, Harvard University, Cambridge, Massachusetts, USA. Supported by DARPA ONISQ program award HR001120C0068. Email: jus065@g.harvard.edu.

§Department of Computer Science, Bocconi University, Milan, Italy. Supported by European Research Council (ERC) award No. 834861. Email: jonathan.shi@unibocconi.it.

Contents

1	Introduction	1
1.1	Diluted k -spin glasses, maximum cut of sparse hypergraphs, and Max- k -XOR	2
1.2	Related work	3
1.3	Our results	5
1.4	Technical overview	8
2	Preliminaries	10
2.1	Spin glass theory	10
2.2	Local algorithms	13
2.3	Overlap-gap properties	16
2.4	Vanishing local neighborhoods of random sparse k -uniform hypergraphs	17
3	Locality and Shared Randomness	17
3.1	Generic p -local algorithms	17
3.2	Locality properties of QAOA for hypergraphs	18
3.3	Shared randomness between runs of a generic local algorithm	19
4	Main Theorems	20
4.1	Obstruction for generic local algorithms on diluted k -spin glasses	20
4.2	Obstructions for generic local algorithms on (k, d) -CSP(f) with coupled OGP	20
4.3	Concentration of objective function values of QAOA	21
5	Concentration Analysis for Generic Local Algorithms	22
5.1	Concentration over the internal randomness of the algorithm	23
5.2	Concentration over randomness of problem instance	27
5.3	Proof of Theorem 5.1	29
5.4	Corollaries of concentration analysis	30
6	Proofs of main theorems	33
6.1	Proof of Theorem 4.2	33
6.2	Proof of Theorem 4.5	35
6.3	Proof of Theorem 4.3	35
7	Strengthened McDiarmid's Inequality for biased distributions	36
7.1	Martingales & concentration	36
7.2	Strengthened McDiarmid's inequality	37
8	Overlap-Gap Property for general case random Max-k-XOR	39
8.1	The signed Max- k -XOR hamiltonian	39
8.2	The Guerra-Toninelli interpolations	39
8.3	Coupled Overlap-Gap Property for general case Max- k -XOR	40

9	Discussion	46
9.1	Which CSPs have an OGP?	46
9.2	Beyond log-depth obstructions for $QAOA_p$?	46
9.3	A Quantum OGP and lifting “classical” obstructions	47
9.4	Message-Passing algorithm for MAX-CUT of <i>all</i> d -regular graphs?	47
A	Proof of Proposition 3.2	54
B	Proof of Theorem 2.20	56

1 Introduction

Recent developments [AAB⁺20, GWZ⁺21, EWL⁺21] of noisy intermediate-scale quantum (NISQ) devices [Pre18] have brought us to the door of near-term quantum computation. As experimentalists can now build programmable quantum simulators up to 256 qubits [EWL⁺21], this motivates an important theoretical question: what computational advantage can such a NISQ device provide?

One of the constraints of NISQ devices is the inability to create high-fidelity global entanglement. This motivates the study of the power of quantum algorithms that are *local*. A leading candidate in this regime of quantum algorithms is the *Quantum Approximate Optimization Algorithm* (QAOA) [FGG14] at shallow depths. While there have been some recent results [Has19, BM21, Mar21] that formally examine the QAOA algorithm at depth $p = 1$ or 2 , very few results exist for super-constant depth QAOA [FGG20a, FGG20b].

Given the imminent quest of demonstrating quantum computational advantage, it is important to clarify for what optimization problems can near-term quantum algorithms (such as *local* quantum algorithms) reliably be expected to demonstrate computational advantage.

We show that local *quantum* algorithms, a large natural class of NISQ algorithms, are obstructed by a geometric property of the solution space known as the *coupled Overlap-Gap Property* [CGP⁺19]. We conjecture that this property is satisfied by most CSPs (Conjecture 9.1). Specific problems known to have this property include the diluted k -spin glass Hamiltonian (equivalent to a max-cut problem on random k -hypergraphs) [CGP⁺19], independent set on random graphs [FGG20a], planted clique [GZ19], and many other problems that so far seem to elude efficient algorithms and be algorithmically hard [GJ21]. In this manuscript, we also demonstrate that the random Max- k -XOR problem has this property (see Section 8).

Critical to our approach is a new definition of local algorithms we term *generic local algorithms* (See Section 3). Previous work relating statistical-physics-derived OGPs to local algorithms leveraged the *factors of i.i.d.* framework for local algorithms, which fails to contain local quantum algorithms, as we demonstrate in Proposition 3.2. Our definition of *generic local algorithms* subsumes local quantum and classical algorithms (see Proposition 3.2 and Proposition 3.3) but still satisfies strong concentration properties (see Theorem 5.3 and Theorem 5.4), allowing obstruction techniques for local classical algorithms [CGP⁺19] to apply to the quantum case. Two of our core technical contributions involve showing that the random MAX- k -XOR problem has a coupled Overlap Gap Property (see Section 8) by extending the techniques of Chen et al [CGP⁺19] and deriving a strengthened version of McDiarmid’s inequality for highly-biased random variables using a martingale argument (see Lemma 7.6).

The rest of the paper is organized as follows: In Section 1.1 we give a brief introduction to the motivating spin glass literature, defining the notion of a diluted k -spin glass; in Section 1.2 we introduce the relevant prior work; in Section 1.3 we state our main theorems (informally); in Section 1.4 we briefly explain the architecture of our proof and compare our techniques with those of Chen et al. [CGP⁺19] and Farhi et al. [FGG20a]; in Section 2 we introduce the necessary mathematical

preliminaries and notation, including a rigorous definition of local classical algorithms, the QAOA algorithm and Overlap-Gap Properties; in [Section 3](#) we introduce the notion of a *generic local algorithm*, how to sample from correlated runs of them, and finally show separation of different families of local algorithms; in [Section 4](#) we state our main theorems formally and give proof sketches; in [Section 5](#) we state and prove multiple concentration of measure statements about *generic local algorithms*; in [Section 6](#) we make the proof showing obstructions against *generic local algorithms* using the same interpolation procedure of Chen et al. [\[CGP⁺19\]](#); in [Section 7](#) we state and prove a strengthened version of *McDiarmid's* inequality; in [Section 8](#) we state and prove an OGP for random Max- k -XOR; in [Section 9](#) we conclude by summarizing our results and mentioning many natural open problems closely related to and/or motivated by our work.

1.1 Diluted k -spin glasses, maximum cut of sparse hypergraphs, and Max- k -XOR

Spin glass theory is a central theoretical framework in statistical physics. The Sherrington-Kirkpatrick model (SK model) [\[SK75\]](#) is one of the most well studied mathematical models in the theory and consists of two variables: spins $\{\sigma_i\}_{i \in [n]}$ and interactions $\{J_{i,j}\}_{i,j \in [n]}$. A spin σ_i takes values in $\{\pm 1\}$ and the interaction $J_{i,j}$ between two spins σ_i, σ_j is a real-valued variable that captures whether the physical system prefers the two spins to be the same ($J_{i,j} > 0$) or different ($J_{i,j} < 0$). The goal is to understand what spin configurations $\sigma \in \{-1, 1\}^n$ maximize the following quantity (a.k.a. Hamiltonian):

$$H(\sigma) = \sum_{i,j} J_{i,j} \sigma_i \sigma_j.$$

The setting is easily generalized to higher order interactions, i.e., $J_{i_1, \dots, i_k}$ acting on k spins, and this is known as the k -spin model. See Panchenko [\[Pan14a\]](#) for a comprehensive survey.

There is a natural correspondence between spin glass theory and combinatorial optimization problems. In a combinatorial optimization problem (e.g., Max-Cut), a variable corresponds to a spin and a constraint corresponds to an interaction. Through this correspondence, the maximization of the above Hamiltonian $H(\sigma)$ serves as a proxy for maximizing the number of satisfied constraints in the combinatorial optimization problem.

A spin glass model additionally specifies a particular distribution on the interactions $\{J_{i,j}\}$ for all $i, j \in [n]$. The quantity of interest is the asymptotic maximum value

$$H^* := \lim_{n \rightarrow \infty} \frac{1}{n} \max_{\sigma} H(\sigma),$$

(a.k.a. the ground state energy density). Also of interest are spin configurations σ with $H(\sigma) \approx H^*$. There are many well-studied spin glass models in physics and various mathematical insights about these have been discovered over the years [\[Con18, DMS17, Sen18, PT04\]](#). For example, for the SK model [\[SK75\]](#) Parisi [\[Par80\]](#) proposed the infamous *Parisi Variational Principle* to capture the exact value of H^* . This was later rigorously proved by Talagrand [\[Tal06\]](#) and again by Panchenko [\[Pan14b\]](#) in greater generality. These successes give hope to design local algorithms that simulate the physical system and output a final configuration as an approximation to the corresponding combinatorial optimization problem.

While traditional spin glass models consider the underlying non-trivial interactions as either lying on a certain physically-realistic graph (e.g., the non-zero $J_{i,j}$ form a 2D-grid) or being a *mean field approximation* (for example, where every $J_{i,j}$ is non-trivial), the applications in combinatorial optimization often require the underlying constraint graphs to be sparse and arbitrary. We use two methods of bridging the gap between the two settings:

- By studying the *diluted k -spin glass model* where one first samples a sparse hypergraph and then assigns non-trivial interactions on top of its hyperedges. Intuitively, approximating the H^* of the diluted k -spin glass corresponds to approximating the maximum cut over random sparse hypergraphs. This correspondence is made more precise in [Section 2.1.1](#).
- Using the techniques of *Gaussian interpolation* [\[GT02\]](#) and *Poisson interpolation* [\[FL03, CGP⁺19\]](#) from statistical physics to relate the behavior of random dense spin glass models to random sparse CSPs. More specifically, we relate the random Max- k -XOR problem to mean-field p -spin glasses ([Section 8](#)).

Spin glass models	Combinatorial optimization problems
Spins $\sigma \in \{-1, 1\}^n$	An assignment to boolean variables
Interactions $\{J_{i_1, \dots, i_k}\}_{i_1, \dots, i_k \in [n]}$	Constraints (i.e., hyperedges)
Hamiltonian $H(\sigma)$	Value of an assignment (i.e., $\text{val}_\Psi(\sigma)$)
Ground state energy H^*	Optimal value (i.e., val_Ψ)
Mean field model (e.g., SK model)	The underlying hypergraph being complete
Diluted spin glass model	The underlying hypergraph being sparse

Table 1: A dictionary between spin glass models and combinatorial optimization problems.

1.2 Related work

Constraint-satisfaction problems & hardness for classical algorithms CSPs (described formally in [Definition 1.4](#)) are a natural class of combinatorial optimization problems that have been studied extensively in theoretical computer science [\[BPS99, Kum92\]](#). Many NP-Complete problems such as k -SAT, k -NAE-SAT, MAX-CUT and k -XOR, can be framed as CSPs. Consequently, unless $\mathcal{P} = \mathcal{NP}$, finding optimal solutions to these problems is infeasible. A natural question then is to understand how well can approximate answers to instances of these problems be constructed by efficient algorithms. Under the now widely believed *Unique Games Conjecture* [\[KV05\]](#), upper bounds on the approximability of CSPs are known [\[KKMO07, Rag08\]](#). These bounds, however, are only worst-case and do not necessarily explicitly demonstrate a family of instances of a CSP that are hard to approximate. Additionally, they remain conditional on a positive resolution to the Unique Games Conjecture, which is still a difficult open problem in the field. In the average-case regime, the goal is to ask how well a *typical* instance of a CSP can be approximated, where the instance is chosen from a “natural” distribution over the set of instances. Perhaps surprisingly, great insight has been drawn about the algorithmic hardness (or lack thereof) about random instances of many CSPs based on work originating in the Statistical Physics community, particularly in Spin Glass

Theory [MP01, FL03, PRTR14]. This was so because the problem of finding spin configurations of particles in many spin glass models that put a system in the ground state could naturally be interpreted as a CSP. Various iterative algorithms were proposed to study the problem of explicitly finding near-ground states of *typical* instances of various spin glass models [YFW⁺03, BMZ05]. It was observed that these algorithms either consistently got better with the number of iterations, or hit a threshold which they could not exceed. To understand this, the work of Achlioptas et al [ART06] studied the solution geometry of the k-SAT problem and found that most good solutions were in well separated clusters. Additionally, most variables in a good solution could only take a single value (i.e., they were “frozen”). This observation was stated as an intuitive reason for the failure of local algorithms on random instances of k-SAT. Gamarnik et al [GS14] made this more formal and precise by showing that *no* classical local algorithm (described formally as *factors of i.i.d.*, see Section 2.2.1) could approximate the MAX-IND-SET problem arbitrarily well on sparse random graphs. Critical to their argument was the fact that *all* (not *most*) nearly-optimal solutions to the problem satisfied the *Overlap Gap Property* - they were in well separated clusters. In various works that followed up, many problems have been shown to have near-optimal solutions conform to this solution geometry and algorithmic hardness for various families of classical algorithms has been established [CGP⁺19, GJW20, GJ21].

Results about QAOA. In their seminal work, Farhi et al. [FGG14] introduced QAOA as a possible way to approximately solve certain hard combinatorial optimization problems. To illustrate the capabilities of QAOA, its performance at $p = 1$ was shown to achieve an approximation ratio of at least .6924 for the Max-Cut problem on triangle-free 3-regular graphs [FGG14]. In a follow up work, Wurtz et al. [WL20] improved this to .7559 for 3-regular graphs with $p = 2$ and made the empirical observation that the bound was tight for graphs with no cycles of length < 7 . Shortly after $QAOA_p$ was proposed, however, a local classical algorithm was designed that outperformed it on these graphs at depth 1 [Has19]. Consequently, because of a flurry of follow up results, QAOA has been shown to be outmatched by local classical algorithms up to depth 2 [Mar21, BM21] for the MAX-CUT problem on d -regular graphs with large girth. In fact, under the widely-believed conjecture in the Spin-Glass Theory community that the SK model does not satisfy the *Overlap Gap Property* [ACZ20], an AMP algorithm was recently proposed that outputs arbitrarily good cuts for large (but constant) degree random regular graphs [AMS21]. However, this result [AMS21] does not *completely* rule out a possibility for quantum advantage (see Section 9.4). To analyze the performance of $QAOA_p$ on a problem that possesses an OGP, Farhi et al. [FGG20a] established that $QAOA_p$ with depth $p \leq \epsilon \log(n)$ could not output independent sets of size better than .854 times the optimal for sparse random graphs. This work suggested that the OGP may broadly prove to be an obstacle for $QAOA_p$ while it is local as much as it does for various classical algorithms. However, MAX-IND-SET is not a (maximum) CSP and, additionally, the prior work [FGG20a] does not give an analysis that generalizes to CSPs. Our work establishes this generalization and also immediately positively resolves the “landscape independence” conjecture of $QAOA_{\epsilon \log(n)}$ proposed by Brandao et al. [BBF⁺18]. This immediately suggests that quantum advantage is unlikely to be found up to this depth for CSPs with an OGP, and we conjecture that *almost all* CSPs will have an OGP (see Section 9.1).

It was shown by Bravyi et al. [BKKT20] that QAOA would not output cuts better than $\frac{5}{6} + O(\frac{1}{\sqrt{d}})$ times the optimal value for some infinite family of d -regular graphs (which happen to be bipartite). This was achieved as a corollary to their proof for a $\log(n)$ -depth version of the NLTS conjecture. Farhi et al. [FGG20b] improved on this via an indistinguishability argument which utilized the fact that local neighborhoods of random d -regular graphs are trees with high probability to then conclude that there are d -regular graphs (specifically random bipartite ones) on which QAOA wouldn't do better than $\frac{1}{2} + O(\frac{1}{\sqrt{d}})$ for sufficiently large n .

Local algorithms for spin glasses. The performance and limitations of various algorithms, such as factors of i.i.d. and message passing algorithms, have been established on different models of spin glasses [CGP⁺19, GJ21, Mon19, EAMS21]. In particular, the literature often provides two kinds of results: An arbitrary approximation to the ground state in the *absence* of an Overlap Gap Property via an appropriate algorithm [EAMS21, Mon19] or a barrier to arbitrary approximation for some family of algorithms in the *presence* of an Overlap Gap Property [CGP⁺19, GJ21]. The first work, to the authors' knowledge, that analyzed the performance of QAOA on a spin glass model was by Farhi et al. [FGGZ19]. In this work, [FGGZ19] provide an analytic expression for the expected value that QAOA_p outputs on typical instances of the SK model, which can be evaluated by a "looping procedure" implemented on a circuit with $O(16^p)$ gates. Numerical results are provided demonstrating evidence that at $p = 11$ this beats the best known SDP-based solver. In this paper, we show that the *Overlap Gap Property* of diluted k -spin glasses poses an obstacle for fixed angle QAOA_p when $p < \epsilon \log(n)$. The generalization to the k -spin mean field model is substantially more challenging to analyze, as in that setting the QAOA_p algorithm is not local even at depth $p = 1$. However, a *coupled* Overlap Gap Property is known to exist for the k -spin mean field model [CGP⁺19].

1.3 Our results

In this work, we show that at shallow-depth the QAOA algorithm cannot output a spin configuration that has Hamiltonian $(1 - \epsilon_0)$ -close to the H^* in a random diluted k -spin glass.

Theorem 1.1 (Obstruction to QAOA over diluted k -spin glass, informal).

For every even $k \geq 4$, there exists $d_0 \in \mathbb{N}$ and the following holds: There exists $\epsilon_0 > 0$ such that if QAOA_p outputs a solution $\sigma \in \{-1, 1\}^n$ with $H(\sigma)$ being $(1 - \epsilon_0)$ -close to the H^ of a random diluted k -spin glass of average degree $d \geq d_0$ with probability at least 0.99, then $p = \Omega(\log n)$.*

The formal version of the theorem is stated in Theorem 4.2. This result can be interpreted as a weak obstruction to logarithmic-depth QAOA in approximating a random diluted k -spin glass, which is equivalent to the random Max- k -XOR problem when all clauses check for odd parity of non-negated variables. We also demonstrate the same result for the general case of random Max- k -XOR.

Theorem 1.2 (Obstruction to QAOA on random Max- k -XOR, informal).

For every even $k \geq 4$, there exists $d_0 \in \mathbb{N}$ and the following holds: There exists $\epsilon_0 > 0$ such that if QAOA_p outputs a solution $\sigma \in \{-1, 1\}^n$ with $H(\sigma)$ being $(1 - \epsilon_0)$ -close to the H^ of a random Max- k -XOR instance of average degree $d \geq d_0$ with probability at least 0.99, then $p = \Omega(\log n)$.*

This is stated formally in [Corollary 4.4](#), and answers a question of [\[FGGZ19\]](#), where the authors ask if QAOA_p would perform well on k -spin generalizations of the SK model, citing Max- k -XOR in particular [\[CGP⁺19\]](#).

In fact, we can prove results stronger in three ways: (i) the same approximation resistance holds for a more general family of algorithms defined as generic local algorithms ([Definition 3.1](#)), (ii) the same approximation resistance holds for a broader family of optimization problems ([Definition 1.4](#)) provided they exhibit a certain solution geometry ([Definition 1.5](#)), and (iii) we show that random Max- k -XOR with negations is one of the optimization problems with this geometry ([Theorem 8.12](#)), whereas previous work only handled it without negations so that all clauses needed to be of odd parity. We begin by informally introducing generic local algorithms, random constraint satisfaction problems (CSPs), and the coupled overlap gap property (OGP).

Generic local algorithms. As traditional notions of local algorithms do not capture QAOA^1 , we generalize the definition to a broader family and call it *generic local algorithms*. A randomized algorithm A on a hypergraph $G = (V, E)$ can be viewed as outputting labels $A(G) \in S^V$ from a label set S (e.g., $S = \{-1, 1\}$). As both A and G are random, $A(G)$ is a set of random variables and the *independence structure* of $A(G)$ captures how local A is. Next, for a hypergraph G and a vertex set $L \subset G$, the p -neighborhood of L is the induced subgraph of G by the vertices that can reach L in p steps.

Definition 1.3 (Generic local algorithms, informal).

Let $p \in \mathbb{N}$ and let S be a finite label set. We say an algorithm A (which takes a hypergraph G as an input) is *generic p -local* if the following hold:

- (Local distribution determination.) For every set of vertices $L \subset V$, the joint marginal distribution of the labels $(A(G)_v)_{v \in L}$ depends only on the union of the p -neighborhoods of $v \in L$ in G .
- (Local independence.) $A(G)_v$ is statistically independent of the joint distribution of $\{A(G)_{v'}\}$ for every v' that is farther than a distance of $2p$ from v .

The main difference between our notion of generic local algorithms and the ones used by previous works is that [Definition 1.3](#) captures the *evolution of correlations*, without assuming any concrete model of randomness. This is crucial in the interpolation step of the proof (see [Section 1.4](#)). See [Definition 3.1](#) for the formal definition.

Definition 1.4 (Random (k, d) -CSP(f)).

A (signed) random (k, d) -CSP(f) instance with a local constraint function $f : \{-1, 1\}^k \rightarrow \{0, 1\}$ is constructed as follows:

1. Choose $r \sim \text{Poisson}(dn/k)$.
2. Sample r clauses of size k by choosing each clause C_i independently as a collection of k variables uniformly at random from $\{x_1, \dots, x_n\}^k$, and, in the case of a signed random CSP, random signs $s_{i,1}$,

¹This is made formal in [Proposition 3.2](#)

$$\dots, s_{i,k} \in \{\pm 1\}.$$

To each clause C_i there are k variables associated: $\{x_{i_1}, \dots, x_{i_k}\}$. A clause is satisfied if there is some assignment to every $x_{i_j} \in \{-1, 1\}$, such that, $f(x_{i_1}, \dots, x_{i_k}) = 1$ (or $f(s_{i,1}x_{i_1}, \dots, s_{i,k}x_{i_k}) = 1$ if signed). The value of an assignment $\sigma \in \{-1, 1\}^n$ is defined as $\text{val}_\Psi(\sigma) := \#\{C_i : f(\sigma_{i_1}, \dots, \sigma_{i_k}) = 1\}$ (or $\#\{C_i : f(s_{i,1}\sigma_{i_1}, \dots, s_{i,k}\sigma_{i_k}) = 1\}$ if signed). The optimal value of Ψ is defined as $\text{val}(\Psi) := \max_\sigma \text{val}_\Psi(\sigma)$.

When unspecified, we will be referring to unsigned CSPs.

We say that a random (un)signed (k, d) -CSP(f) satisfies a *coupled overlap-gap property (OGP)* if, given two instances Ψ, Ψ' constructed so that they share a random t -fraction of clauses with the remaining $(1 - t)$ -fraction chosen independently, any two “good” solutions σ of Ψ and σ' of Ψ' are either very similar or dissimilar.

Definition 1.5 (Coupled OGP, informal).

A signed or unsigned (k, d) -CSP(f) satisfies a coupled OGP if there exists $\epsilon_0 > 0$ and $0 < a < b < 1$ such that the following hold for every $t \in [0, 1]$: Given two (k, d) -CSP(f) instances Ψ, Ψ' constructed so that they share a random t -fraction of their clauses and have the remaining $(1 - t)$ -fraction of clauses chosen independently and uniformly at random, then for every $0 < \epsilon < \epsilon_0$, the overlap between any $(1 - \epsilon)$ -optimal solution σ of Ψ and σ' of Ψ' satisfies

$$\frac{1}{n} \langle \sigma, \sigma' \rangle \notin [a, b]$$

with high probability.

A formal definition of the property above is provided in [Theorem 2.19](#), and the formal definition of the interpolation used to create two “ t -coupled” instances is given in [Definition 2.18](#). Note that an instance of a (k, d) -CSP(f) can be thought of as a random sparse k -hypergraph, and this is made more precise in the proof of [Theorem 4.5](#).

Now, we are able to state the most general form of our main result.

Theorem 1.6 (Obstruction to generic local algorithms given coupled OGP, informal).

For every $k \geq 2$, and a constraint function $f : \{-1, 1\}^k \rightarrow \{0, 1\}$, suppose there exists d_0 where a random signed or unsigned (k, d) -CSP(f) satisfies the coupled OGP (i.e., [Definition 1.5](#)) for every $d \geq d_0$ and $p(n)$ is such that it satisfies the requirements of [Theorem 5.3](#) and [Theorem 5.4](#), then the following holds: There exists $\epsilon_0 > 0$ such that a generic $p(n)$ -local algorithm cannot output a solution that is better than $(1 - \epsilon_0)$ -optimal with high probability.

The formal version of the above theorem is stated in [Theorem 4.3](#). The theorem effectively obstructs any algorithm that makes assignments for variables by looking at $o(n)$ sized local neighborhoods irrespective of how these decisions are made and what kind of randomness is used, provided the problem exhibits a coupled OGP.

Confirmation of landscape independence. As a consequence of our proof techniques, we also confirm a prediction of Brandao et al. [[BBF⁺18](#)] in the $\Theta(\log n)$ -depth regime for QAOA by showing that the output values of QAOA on a random (k, d) -CSP(f) instance (with depth p as stated in

[Lemma 2.20](#)) concentrate very heavily around the expected value. Once again, the expectation here is with respect to the input distribution as well as the internal randomness of the algorithm.

Theorem 1.7 (Confirmation of landscape independence, informal).

Given a random instance Ψ of a (k, d) -CSP(f) chosen as stated in [Definition 1.4](#), and a QAOA _{p} circuit with depth $p < g(d, k) \log(n)$ for some function g , the solution σ output by QAOA _{p} with value $\text{val}_\Psi(\sigma)$ concentrates as,

$$\Pr[|\text{val}_\Psi(\sigma) - \mathbb{E}[\text{val}_\Psi(\sigma)]| \geq \delta n] \leq o_n(1),$$

for every $\delta > 0$ and the probability taken over both the input distribution and internal randomness of the algorithm.

The theorem above is made formal in [Theorem 4.5](#), and the proof follows by encoding a random (k, d) -CSP(f) instance in a random sparse k -hypergraph and then applying [Corollary 5.13](#) with the local function on every hyperedge being set to f .

Discussion and open problems. Our results reveal that a coupled OGP is tightly related to the obstruction of QAOA. This motivates many open problems that are either inspired by this work or are closely related to it, and these are discussed in greater detail in [Section 9](#).

1.4 Technical overview

Our proofs for the main theorems follow the analysis framework of [\[CGP⁺19\]](#), which shows the approximation resistance of random diluted k -spin glasses to a weaker² class of classical algorithms called *factors of i.i.d. local algorithms*. We start with briefly giving an overview of their proof and pointing out where their analysis does not extend to QAOA. See also [Figure 1](#) for a pictorial overview.

Chen et al. [\[CGP⁺19\]](#) analysis. They establish a coupled overlap-gap property (OGP) for diluted k -spin glasses ([Theorem 2.19](#)). The property says that for two “coupled” random instances and any nearly optimal solutions $\sigma_1, \sigma_2 \in \{-1, 1\}^n$ of these, the solutions either have large or small overlap on the assignment values to the variables, i.e., there exists an interval $0 < a < b < 1$ such that $\langle \sigma_1, \sigma_2 \rangle / n \notin [a, b]$. The coupled OGP holds over an interpolation of a pair of hypergraphs $\{(G_1(t), G_2(t))\}_{t \in [0, 1]}$ with the following three properties: for every $t \in [0, 1]$, denote $\sigma_1(t)$ and $\sigma_2(t)$ as the outputs of a factors i.i.d. algorithm on inputs $G_1(t)$ and $G_2(t)$ respectively. (i) when $t = 0$, $(G_1(0), G_2(0))$ are independent random hypergraphs and $\langle \sigma_1(0), \sigma_2(0) \rangle / n < a$ with high probability; (ii) when $t = 1$, $G_1(1) = G_2(1)$ are the same random hypergraph and $\langle \sigma_1(1), \sigma_2(1) \rangle / n = 1$ with high probability; (iii) for each $t \in [0, 1]$, the correlation $\langle \sigma_1(t), \sigma_2(t) \rangle / n$ between the two solutions is highly concentrated (with respect to the randomness of $G_1(t), G_2(t)$ and the algorithm) to a value $R(t)$, and $R(t)$ is a continuous function of t . This contradicts the OGP if the solutions are nearly optimal and hence no such factors of i.i.d. algorithm can exist. Note that it is also important to

²In particular, QAOA is not captured by factors of i.i.d. and we show a separation in [Proposition 3.2](#). Refer to [Appendix A](#) for more details.

assert that the hamming weight and the objective function values output by the algorithm also concentrate.

Our analysis. The key part of the Chen et al. [CGP⁺19] proof that does not work for QAOA is item (iii) of step 2. Specifically, QAOA is not a factors of i.i.d. local algorithm and hence their concentration analysis on the correlation between solutions to coupled instances does not apply. Intuitively, this is because local quantum circuits can induce entanglement between qubits in a local neighborhood which cannot be explained by a local hidden variable theory [Bel64]. We overcome this issue by first generalizing the notion of factors of i.i.d. algorithms to what we call *generic local algorithms* (Definition 3.1).

To establish concentration of overlap for generic local algorithms, the challenge lies in how to capture the local correlations of $G_1(t)$ and $G_2(t)$. We achieve this by defining a new notion of a random vector being *locally independent* (Definition 5.2). The locally independent structure enables us to show concentration on a fixed instance over multiple runs of the generic local algorithm with respect to its internal randomness (Theorem 5.3). Finally, to establish concentration between a pair of correlated instances ($G_1(t)$ and $G_2(t)$), we strengthen McDiarmid’s inequality for biased distributions (Lemma 7.6) and this allows the concentration analysis of the correlation function $R(t)$ to pull through (Theorem 5.1). We complete the analysis by showing that the hamming weight and objective function values output by a generic p -local algorithm also concentrate (Corollary 5.10, Corollary 5.13). In fact, we show this for a broader class of problems (Theorem 4.5).




Figure 1: Overview of the proof ideas. (a) The coupled overlap-gap property (OGP) by [CGP⁺19]. Pictorially, the property guarantees that the nearly optimal solutions of a pair of independent instances form multiple disjoint clouds in the solution space $\{-1, 1\}^n$. (b) Interpolation of a pair of diluted k -spin glass instances. $\{(G_1(t), G_2(t))\}_{t \in [0, 1]}$ describes the interpolation from totally independent (i.e., $t = 0$) to totally correlated (i.e., $t = 1$). $R : [0, 1] \rightarrow [0, 1]$ is the correlation function of a local algorithm (e.g., QAOA) on this interpolation. The coupled OGP prevents $R(t)$ to take values in $[a, b]$ (i.e., the grey area). (c) To contradict the coupled OGP, we would like to show that $R(t)$ is a continuous function and this requires showing that the correlation between the outputs of $G_1(t)$ and that of $G_2(t)$ is concentrated. This is the main challenging step in the proof.

Comparison with Chen et al. [CGP⁺19] and Farhi et al. [FGG20a]. We augment the techniques of [FGG20a] to handle a coupled OGP over a continuous interpolation, as opposed to the coupled

OGP in [FGG20a] which is over a fixed discrete interpolation. The advantage of this is to enable the use of a broader family of coupled OGP’s provable using statistical mechanics methods, whereas the coupled OGP of [FGG20a] requires reasoning about explicit sequences of instances in a way that does not clearly generalize from their independent set analysis to the setting of CSPs. Our statements additionally are more general and show much stronger concentration than those of [CGP⁺19] which is necessary to demonstrate that polynomially many runs of the algorithm will (with high probability) not succeed. We also show a locality bound of $\log n$ instead of constant, requiring different techniques for analyzing locality than those used in [CGP⁺19], who study only regimes where all neighborhoods are locally isomorphic to trees. Finally, we demonstrate the coupled OGP (and therefore obstruct generic local algorithms) for general-case random Max- k -XOR, rather than the case where all clauses require odd parity of their associated variables, without negations.

2 Preliminaries

In this paper, we adopt the following conventions on notations in a CSP (in a spin glass model). n denotes the number of variables (the number particles); k denotes the arity of a constraint (the number of particles involved in an interaction) and $k \geq 4$ throughout the paper; m denotes the number of hyperedges (the total number of non-trivial interactions); d denotes the degree of a variable (the number of interactions a particle is involved in on average).

The rest of this section is organized as follows. We first recall some elementary definitions and results from spin glass theory in Section 2.1. Then, we formally define local algorithms such as factors of i.i.d. and QAOA in Section 2.2. We then provide the complete definition of the OGP and coupled OGP as well as some relevant theorems in Section 2.3. Finally, in Section 2.4 we end with a statement which states that sufficiently local neighborhoods of sparse random hypergraphs see a vanishing fraction of the total hypergraph.

2.1 Spin glass theory

As introduced in Section 1.1, a spin-glass model is specified by a collection of interactions $\{J_{i_1, \dots, i_k}\}$ on n particles. Physicists are interested in studying Hamiltonians of the form,

$$H(\sigma) = \sum_{i_1, \dots, i_k} J_{i_1, \dots, i_k} \sigma_{i_1} \cdots \sigma_{i_k},$$

where $\sigma \in \{-1, 1\}^n$ is the spin configuration. Specifically, it is of importance to understand the *ground state energy*, i.e., $H^* := H(\sigma)$, as well as the spin configurations σ that have energy $H(\sigma)$ close to H^* . Note that this naturally connects spin glass theory to combinatorial optimization problems where $\{J_{i_1, \dots, i_k}\}_{i_1, \dots, i_k \in [n]}$ specifies the input constraints, σ corresponds to the variables, and $H(\sigma)$ is the objective function.

In condensed matter physics, it is critical to understand the average-case setting and, therefore, the interactions $\{J_{i_1, \dots, i_k}\}_{i_1, \dots, i_k \in [n]}$ are sampled from a certain distribution. We now introduce two

common spin glass models: diluted k -spin glasses (Section 2.1.1) and the k -spin mean field model (Section 2.1.2).

2.1.1 Diluted k -spin glasses

In the diluted k -spin glass model, the interactions are sampled from a random sparse k -uniform hypergraph defined as follows.

Definition 2.1 (Hypergraphs on n Vertices).

A hypergraph G on n vertices with m hyperedges is characterized by its set of labelled vertices $V = \{1, \dots, n\}$ and hyperedges $E = \{e_1, \dots, e_m\}$, where every hyperedge $e_i = (v_{i,1}, \dots, v_{i,k}) \in E$ is an ordered k_i -tuple in V , and $k_i \in \mathbb{N}$, $\forall i \in [m]$.

We restrict our attention to sparse instances of such hypergraphs, which amounts to asserting that the number of hyperedges is $m = nd/k = O(n)$. Additionally, we also restrict to the case that our hyperedges are k -uniform, that is, each of them contains k -vertices. For the rest of the paper, we will always assume that $d = O(1)$, $k \geq 4$ and k is even, and that $\mathcal{H}_{n,d,k}$ denotes the set of all such k -uniform hypergraphs over n vertices with nd/k hyperedges.

Definition 2.2 (Random Sparse k -Uniform Hypergraphs).

A hypergraph $G \sim \mathcal{H}_{n,d,k}$ is chosen by first choosing the number of edges $m = |E| \sim \text{Poisson}(dn/k)$, and then choosing hyperedges $e_1, \dots, e_m$ i.i.d. uniformly at random from the set $[n]^k$ of all vertex k -tuples.

Let G be a sparse k -uniform hypergraph, the Hamiltonian of the corresponding diluted k -spin glass is

$$H^G(\sigma) = - \sum_{i=1}^m \prod_{j=1}^k \sigma_{v_{i,j}}, \quad (2.3)$$

where $\sigma_{v_{i,j}}$ denotes the spin of the j -th vertex in the i -th hyperedge. Note that maximizing this Hamiltonian corresponds to finding a configuration σ such that $H^G(\sigma)$ is maximized.

Correspondence to Max- k -XOR. Maximizing the Hamiltonian in a diluted k -spin glass is equivalent to maximizing the number of satisfying constraints in a certain instance Ψ of Max- k -XOR. Recall that a Max- k -XOR instance consists of constraints of the form $x_{i_1} \oplus \dots \oplus x_{i_k} = b_i$. Let $x \in \{0, 1\}^n$ be a boolean assignment, the value of x on Ψ is then defined as $\text{val}_\Psi(x) := \sum_i b_i \oplus x_{i_1} \oplus \dots \oplus x_{i_k}$. Let G be a sparse k -uniform hypergraph, we associate it with a Max- k -XOR instance Ψ^G with constraint $x_{v_{i,1}} \oplus \dots \oplus x_{v_{i,k}} = 0$ for every $i = 1, \dots, m$. Finally, we associate a spin configuration $\sigma \in \{-1, 1\}^n$ to a boolean assignment $x \in \{0, 1\}^n$ by sending $-1 \mapsto 1$ and $1 \mapsto 0$. Thus, the product of spins on a hyperedge e is mapped to the parity of the corresponding boolean variables:

$$\prod_{j=1}^k \sigma_{v_{i,j}} = (-1)^{\bigoplus_{j=1}^k x_{v_{i,j}}}$$

for each $i = 1, \dots, m$. Moreover, the Hamiltonian and the CSP value have the following correspondence.

$$H^G(\sigma) = - \sum_{i=1}^m \prod_{j=1}^k \sigma_{v_{i,j}} = \sum_{i=1}^m \left(2 \cdot (-1)^{\bigoplus_{j=1}^k x_{v_{i,j}}} - 1 \right) = 2\text{val}_{\Psi_G}(x) - m.$$

That is, maximizing that Hamiltonian H^G is equivalent to maximizing the value val_{Ψ_G} . As a remark, note that when $k = 2$, Max- k -XOR becomes Max-Cut. As a *signed* extension of the diluted k -spin glass hamiltonian, one can define a hamiltonian for the random Max- k -XOR problem.

Definition 2.4 (Random Max- k -XOR).

Sample a hypergraph $G \sim \mathcal{H}_{n,d,k}$ by the same procedure mentioned in [Definition 2.2](#). The hamiltonian corresponding to the random Max- k -XOR instance Ψ_G generated by this hypergraph G is

$$H_{\text{signed}}^G = - \sum_{i=1}^m \prod_{j=1}^k p_{ij} \sigma_{v_{ij}}, \quad (2.5)$$

where every $p_{ij} \sim \{\pm 1\}$ is an i.i.d. Rademacher random variable.

Typical behavior. It is important to understand the typical value of H^G when G is a random sparse graph. For example, the following quantity

$$M(k, d) := \lim_{n \rightarrow \infty} \mathbb{E}_{G \sim \mathcal{H}_{n,d,k}} \left[\max_{\sigma \in \{-1,1\}^n} \frac{H^G(\sigma)}{n} \right], \quad (2.6)$$

is a well-defined limit whose existence is inferred from arguments similar to those presented in [\[DS04\]](#). Furthermore, by standard concentration arguments, the ground state energy concentrates heavily around $M(k, d)$. The exact computation of the value of $M(k, d)$ is beyond the scope of this paper. However, as mentioned in [Section 2.1.2](#), the value can be related to the free-energy of a typical instance of the k -spin mean field Hamiltonian in the large- d regime.

2.1.2 k -spin mean field model

The k -spin mean field model is a special case of the infinite-range model with each interaction $J_{i_1, \dots, i_k} = g_{i_1, \dots, i_k} / \sqrt{n^{k-1}}$ where $g_{i_1, \dots, i_k}$ are i.i.d. standard Gaussian random variables. Just as in the case of diluted k -spin glasses, spin configurations that maximize the Hamiltonian are of particular interest. Specifically, we are interested in typical ground state configurations in the thermodynamic limit ($n \rightarrow \infty$). The optimal (normalized) value of the ground state is characterized by the following term,

$$P(k) := \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}_{g_{i_1, \dots, i_k}} \left[\max_{\sigma \in \{-1,1\}^n} H(\sigma) \right], \quad (2.7)$$

where $P(k)$ denotes the famous Parisi constant. In a sequence of recent works [\[DMS17, Sen18\]](#) this limit was precisely related to the limit of the ground state energy of diluted k -spin glasses in the large degree limit as,

$$M(k, d) \xrightarrow{d \rightarrow \infty} \sqrt{\frac{d}{k}} \cdot P(k) + o\left(\sqrt{d}\right). \quad (2.8)$$

2.2 Local algorithms

A local algorithm assigns a (random) label to each vertex v independently at the beginning and then updates it based on the labels of a small neighborhood of v . Intuitively, the labels associated to the vertices form a stochastic process and in the end the local algorithm assigns a value to each vertex according to its final label. In [Section 2.2.1](#) we introduce factors of i.i.d. algorithms [[GS14](#), [CGP⁺19](#)]. These algorithms parameterize a family of local algorithms that capture most common classical local algorithms. We then introduce the QAOA in [Section 2.2.2](#).

2.2.1 Factors of i.i.d. algorithms

A local algorithm takes an input (hyper-) graph G and a label set S , runs a stochastic process $\{X^G(t)\}_t$ that associates to each vertex v a label $X_v^G(t) \in S$ at time t , and outputs the assignment σ_v to each vertex v according to its final label. While there is a huge design space for local algorithms, a factors of i.i.d. algorithm of radius p has the following restrictions: (i) the initial label for each vertex v is set to be an i.i.d. set of random variables $X_v^G(0)$. (ii) For each vertex v , the assignment σ_v is a random variable that only depends on the labels from a p -neighborhood of v . (iii) The assignment function for each vertex is the same. Common local algorithms such as Glauber dynamics and Belief Propagation are examples of factors of i.i.d. algorithms.

To be more concrete, let us start with a formal definition of the p -neighborhood of a vertex in a hypergraph, which is a generalization from the p -neighborhood of a graph by considering two vertices v and w to be adjacent if they belong to the same hyperedge e .

Definition 2.9 (p -neighborhood and hypergraphs with radius p).

Let G be a hypergraph, $v \in V(G)$, and $p \in \mathbb{N}$. The p -neighborhood of v is defined as

$$B_G(v, p) := \{w \in V(G) \mid w \text{ is } p \text{ hyperedges away from } v\}.$$

Let G be a hypergraph, $v \in V(G)$, and $p \in \mathbb{N}$. We say (G, v) has radius p if $B_p(G_v) = G$. Further, let $k \in \mathbb{N}$, we define

$$\mathcal{G}_p := \{(G, v) \mid (G, v) \text{ has radius } p \text{ and } G \text{ is connected, finite, and } k\text{-uniform}\}$$

be the collections of hypergraphs with radius at most p .

Next, to capture the fact that local algorithms assign the value of a vertex v by only looking at a p -neighborhood, it is natural to define an equivalent classes of local induced subgraphs rooted at v as follows.

Definition 2.10 (Rooted-isomorphic graphs).

Let G_1, G_2 be two hypergraphs and $v \in V(G_1) \cap V(G_2)$. We say G_1 and G_2 are rooted-isomorphic at v , denoted as $G_1 \cong_v G_2$, if there exists a hypergraph isomorphism $\phi : V(G_1) \rightarrow V(G_2)$ such that $\phi(v) = v$. Similarly, let $L \subseteq V(G_1) \cap V(G_2)$, we say $G_1 \cong_L G_2$ if there exists a hypergraph isomorphism $\phi : V(G_1) \rightarrow V(G_2)$ with $\phi(v) = v$ for all $v \in L$.

In the future usage of [Definition 2.10](#), we think of G_1 and G_2 as some neighborhoods. Intuitively, when the neighborhood of v_1 and v_2 are rooted-isomorphic, then the local algorithm will give the same output to them.

The last notion of local algorithms to capture is the assigning process f from the labels of a p -neighborhood to a value. In particular, a local algorithm should produce the same output value for v_1 and v_2 when the induced subgraphs of their p -neighborhood are rooted-isomorphic. For simplicity, we focus on the case where the label set $S = [0, 1]$.

Definition 2.11 (Factor of radius p , [\[CGP⁺19, Section 2\]](#)).

Let $p \in \mathbb{N}$. We define the collection of all $[0, 1]$ -labelled hypergraphs of radius at most p as

$$\Lambda_p := \left\{ (G, v, X) \mid (G, v) \in \mathcal{G}_p \text{ and } X \in [0, 1]^{V(G)} \right\}$$

We say $(G_1, v_1, X_1), (G_2, v_2, X_2) \in \Lambda_r$ are isomorphic if there exists a hypergraph isomorphism $\phi : V(G_1) \rightarrow V(G_2)$ such that (i) $\phi(v_1) = v_2$ and (ii) $X_2 \circ \phi = X_1$.

Finally, we say $f : \Lambda_p \rightarrow \{-1, 1\}$ is a factor of radius p function if (i) it is measurable and (ii) $f(G_1, v_1, X_1) = f(G_2, v_2, X_2)$ for every isomorphic $(G_1, v_1, X_1), (G_2, v_2, X_2) \in \Lambda_p$.

Now, we are ready to define factors of i.i.d. algorithms. Intuitively, the output distribution of a factors of i.i.d. algorithm with radius p on a vertex v is determined by the p -neighborhood of v .

Definition 2.12 (Factors of i.i.d., [\[CGP⁺19, Section 2\]](#)).

Let $k, p \in \mathbb{N}$. A factors of i.i.d. algorithm A with radius p is associated with a factor of radius p function f with the following property. On input a k -uniform hyper graph G , the algorithm A samples a random labeling $X = \{X(v)\}_{v \in V(G)}$ where $X(v)$'s are i.i.d. uniform random variables on $[0, 1]$. The output of A is $\sigma \in \{-1, 1\}^{V(G)}$ where

$$\sigma_v := f(B_p(G, v), v, \{X(w)\}_{w \in B_p(G, v)})$$

for each $v \in V(G)$.

2.2.2 The QAOA algorithm

The algorithm. The QAOA algorithm was proposed by Farhi et al. [\[FGG14\]](#) as a way to approximately solve hard combinatorial optimization problems. The QAOA algorithm works by applying, in alternation, weighted rotations in the X -basis to introduce mixing over the uncertainty in the solution space and weighted cost $e^{-i\beta_j H_c(G)}$ unitaries to introduce correlation spreading encoded by the Hamiltonian of the desired cost-function to maximize. This weighting is accomplished by giving some assignment of weights to weight vectors $\hat{\beta} = (\beta_1, \dots, \beta_p)$ and $\hat{\gamma} = (\gamma_1, \dots, \gamma_p)$, and then running a classical optimizer to help find the ones that maximize the output of QAOA. Various methods, including efficient heuristics, to optimize these angles are studied in the literature [\[BBF⁺18, ZWC⁺20\]](#).

The QAOA circuit parametrized by angle vectors $\hat{\gamma}$ and $\hat{\beta}$ looks as follows,

$$U_p(\hat{\beta}, \hat{\gamma}) = \prod_{j=1}^p e^{-i\beta_j \sum_{k=1}^n X_k} e^{-i\gamma_j H_C(G)}.$$

Typically, the initial state on which the circuit is applied is a *symmetric product state*, most notably $|0\rangle^{\otimes n}$ or $|+\rangle^{\otimes n}$. The expected value that QAOA outputs after applying the circuit on some initial state $|\psi_0\rangle$ is,

$$\langle \psi_0 | U^\dagger H_C(G) U | \psi_0 \rangle.$$

It is the expectation value above that is optimized (by maximizing) for various choices of $\hat{\beta}$ and $\hat{\gamma}$ under a classical optimizer, and the solution corresponding to this solution comes from a measurement in the Z-basis of the state $U_p(\hat{\beta}, \hat{\gamma}) |\psi_0\rangle$. We will notate by $\text{QAOA}(\hat{\beta}, \hat{\gamma})$ a QAOA circuit of depth $2p$ with angle parameters $\hat{\beta}$ and $\hat{\gamma}$. In our regime, we will work with *any* collection of *fixed* angles $(\hat{\beta}, \hat{\gamma})$. The fixed angles regime is necessary when reasoning about the concentration of overlaps of the solutions produced by a p -local algorithm on coupled instances with shared randomness. If the angles of QAOA vary between the coupled instances, then we cannot assert that the coupled instances will share randomness when labeling vertices with identical neighborhoods.

The diluted k -spin glass Hamiltonian and QAOA. We can rewrite $H_{k,d,n}$ as a Hamiltonian for a quantum system by replacing σ_i with the Pauli Z matrix. This yields a k -local Hamiltonian that the QAOA ansatz tries to maximize. The Hamiltonian is,

$$\hat{H}_{k,d,n}^G(\sigma) = - \sum_{(v_1, \dots, v_k) \in E(G)} \bigotimes_{i=1}^k \sigma_z(v_i), \quad (2.13)$$

where $\sigma_z(v_i)$ is a 2×2 Pauli Z matrix for the i -th vertex (qubit) in the hypergraph G . We want to maximize the following expectation value,

$$\max_{\alpha_j, \beta_j, j \in [p]} \langle \psi_0 | U^\dagger H_{k,d,n}(G) U | \psi_0 \rangle = \max_{\alpha_j, \beta_j, j \in [p]} - \sum_{(v_1, \dots, v_k) \in E(G)} \left(\langle \psi_0 | U^\dagger \left(\bigotimes_{i=1}^k \sigma_z(v_i) \right) U | \psi_0 \rangle \right). \quad (2.14)$$

QAOA at shallow depth. Note that the only "spreading" of correlation is introduced by the operator $e^{-i\beta_j H_C(G)}$ which is applied only p times. The hamiltonian in consideration is k -local, and therefore, after p operations a qubit i will interact with no more than

$$|B_G(v_i, p)| \leq \left((k-1) \cdot \max_{i \in [n]} |\Pi_i(E(HG))| \right)^p$$

vertices, where,

$$\Pi_i(E(G)) = \{e \in E(G) \mid e \text{ is a hyperedge that contains } v_i\}. \quad (2.15)$$

Proposition 3.3 makes a precise statement about the locality of QAOA with fixed angles. Bounding the size of $\max_{i \in [n]} |\Pi_i(E(HG))|$ is the main subject of [Lemma 2.20](#), which parameterizes p appropriately as a function of the number of vertices n in the graph (logarithmic) as well as the parameters of k and d so that this size is $o(n)$. This is sufficiently small for the purposes of our obstruction theorem.

2.3 Overlap-gap properties

We now state the OGP as it holds for the diluted k -spin glass model in both uncoupled and coupled form. To do so, we begin by introducing the notion of an overlap between two spin-configurations σ_1 and σ_2 , which is equivalent to the number of spins that are the same in both configurations subtracted by the number of different spins, normalized by the number of particles in the system. Formally,

Definition 2.16 (Overlap between spin configuration vectors).

Given any two vectors $\sigma_1, \sigma_2 \in \{-1, 1\}^n$, the overlap between them is defined as,

$$R(\sigma_1, \sigma_2) = \frac{1}{n} \langle \sigma_1, \sigma_2 \rangle = \frac{1}{n} \sum_{i \in [n]} (\sigma_1)_i (\sigma_2)_i.$$

We first state the OGP for diluted k -spin glasses about the overlap gaps in a *single instance*.

Theorem 2.17 (OGP for Diluted k -Spin Glasses, [CGP⁺19, Theorem 2]).

For every even $k \geq 4$, there exists an interval $0 < a < b < 1$ and parameters $d_0 > 0$, $0 < \eta_0 < P(k)$ and $n_0 > 1$, such that, for $d \geq d_0$, $n \geq n_0$ and $L = L(\eta_0, d)$, with probability at least $1 - Le^{-n/L}$ over the random hypergraph $G \sim \mathcal{H}_{n,d,k}$, whenever two spins σ_1, σ_2 satisfy

$$\frac{H^G(\sigma_i)}{n} \geq M(k, d) \left(1 - \frac{\eta_0}{P(k)} \right),$$

then also, $|R(\sigma_1, \sigma_2)| \notin (a, b)$.

A more general version of the OGP excludes, with high probability, a certain range of overlaps between any two solutions of two different instances jointly drawn from a coupled random process. We first introduce this process, and then state the *coupled* version of the OGP as proven in [CGP⁺19].

Definition 2.18 (Coupled Interpolation, [CGP⁺19, Section 3.2]).

The coupled interpolation $\mathcal{H}_{d,k,n,t}$ generates a coupled pair of hypergraphs $(G_1, G_2) \sim \mathcal{H}_{d,k,n,t}$ as follows:

1. First, a random number is sampled from $\text{Poisson}(tdn/k)$, and that number of random k -hyperedges are uniformly drawn from the set $[n]^k$ and put into a set E .
2. Then, two more random numbers are independently sampled from $\text{Poisson}((1-t)dn/k)$, and those numbers of random k -hyperedges are independently drawn from $[n]^k$ to form the sets E_1 and E_2 respectively.
3. Lastly, the two hypergraphs are constructed as $G_1 = (V, E \cup E_1)$ and $G_2 = (V, E \cup E_2)$.

Theorem 2.19 (OGP for Coupled Diluted k -Spin Glasses, [CGP⁺19, Theorem 5]).

For every even $k \geq 4$, there exists an interval $0 < a < b < 1$ and parameters $d_0 > 0$, $0 < \eta_0 < P(k)$ and $n_0 > 1$, such that, for any $t \in [0, 1]$, $d \geq d_0$, $n \geq n_0$ and constant $L = L(\eta_0, d)$, with probability at least $1 - Le^{-n/L}$ over the hypergraph pair $(G_1, G_2) \sim \mathcal{H}_{n,d,k,t}$, whenever two spins σ_1, σ_2 satisfy

$$\frac{H^{G_i}(\sigma_i)}{n} \geq M(k, d) \left(1 - \frac{\eta_0}{P(k)} \right),$$

then their overlap satisfies $|R(\sigma_1, \sigma_2)| \notin [a, b]$.

We also provide a corresponding coupled OGP for random Max- k -XOR in [Theorem 8.12](#).

2.4 Vanishing local neighborhoods of random sparse k -uniform hypergraphs

We state a bound on sufficiently local neighborhoods of random sparse k -uniform hypergraphs.

Lemma 2.20 (Vanishing local neighborhoods of random sparse k -uniform hypergraphs).

Let $k \geq 2$ and $d \geq 2$ and $\tau \in (0, 1)$. Then there exists $a > 0$ and $0 < A < 1$, such that, for n large enough and p satisfying

$$2p + 1 \leq \frac{(1 - \tau) \log n}{\log\left(\frac{d(k-1)}{\ln 2}\right)},$$

the following are true:

$$\Pr_{G \sim \mathcal{H}_{n,d,k}} [\max_i B_G(v_i, 2p) \geq n^A] \leq e^{-n^a},$$

and

$$\Pr_{G \sim \mathcal{H}_{n,d,k}} [\max_i B_G(v_i, p) \geq n^{\frac{A}{2}}] \leq e^{-n^{\frac{a}{2}}}.$$

Intuitively, the above lemma says that the local neighborhood of each vertex is vanishingly small with high probability. To prove [Lemma 2.20](#), we utilize a modified version of the proof of Farhi et al. [[FGG20a](#), Neighborhood Size Theorem] to handle the case of sparse random hypergraphs and we defer the complete proof to [Appendix B](#).

3 Locality and Shared Randomness

3.1 Generic p -local algorithms

We introduce a concept of “local random algorithm” which will allow for different runs of the same local algorithm to “share their randomness”, even when run on mostly-different instances. Later we will demonstrate that QAOA is a local algorithm under this definition.

Definition 3.1 (Generic local algorithms).

We consider randomized algorithms on hypergraphs whose output $A(G) \in S^V$ assigns a label from some set S to each vertex in V . Such an algorithm is generic p -local if the following hold.

- (Local distribution determination.) For every set of vertices $L \subset V$, the joint marginal distribution of its labels $(A(G)_v)_{v \in L}$ is identical to the joint marginal distribution of $(A(G')_v)_{v \in L}$ whenever $\bigcup_{v \in L} B_G(v, p) \cong_L \bigcup_{v \in L} B_{G'}(v, p)$, and,
- (Local independence.) $A(G)_v$ is statistically independent of the joint distribution of $A(G)_{v'}$ over all $v' \notin B_G(v, 2p)$.

Consequently, it will be possible to sample $A(G)_v$ without even knowing what the hypergraph looks like beyond a distance of p away from v .

This definition is more general than the factors of i.i.d. concept used in probability theory [GS14, CGP⁺19]. Our definition, for instance, encompasses local quantum circuits whereas factors of i.i.d. algorithms satisfy Bell's inequalities and do not capture quantum mechanics.

Proposition 3.2 (Generic local strictly generalizes factors of i.i.d.).

There exists a generic 1-local algorithm as defined in Section 3 that is not a 1-local factors of i.i.d. algorithm as defined in Section 2.2.1.

A proof of this proposition is provided in Appendix A, and consists of setting up a Bell's inequality experiment within the framework of a generic 1-local algorithm.

3.2 Locality properties of QAOA for hypergraphs

We show that any QAOA circuit of depth p with some *fixed* angle parameters $(\hat{\beta}, \hat{\gamma})$ is a p -local algorithm. This allows us to describe a process to sample outputs of this circuit when it is run on two different input hypergraphs.

Proposition 3.3.

For every $p > 0$, angle vectors $\hat{\beta}$ and $\hat{\gamma}$, $\text{QAOA}_p(\hat{\beta}, \hat{\gamma})$ is generic p -local under Definition 3.1.

Proof. To see this, consider the structure of QAOA: we start with a product state $|\psi_0\rangle$ where each qubit corresponds to a vertex in the hypergraph, apply the unitary transformation $U = U_p(\hat{\beta}, \hat{\gamma})$ to the state, and then measure each vertex v in the computational basis with the Pauli-Z operator $\sigma_z(v)$. Equally valid and equivalent is the Heisenberg picture interpretation of this process, where we keep the product state $|\psi_0\rangle$ fixed but transform the measurements according to the reversed unitary transformation $U^\dagger$, so that we end up taking the measurements $U^\dagger \sigma_z(v) U$ on the fixed initial state.

Because the $\sigma_z(v)$ operators all commute with each other, their unitarily transformed versions $U^\dagger \sigma_z(v) U$ also mutually commute, and the measurements can be taken in any order without any change in results. Let $M(v) = U^\dagger \sigma_z(v) U$ and $M(L) = \{U^\dagger \sigma_z(u) U \mid u \in L\}$.

To show that QAOA satisfies the first property of generic p -local algorithms, we need to show that the marginal distribution of its assignments to any set $L' \subseteq V$ of vertices depends only on the union of the p -distance neighborhoods of L' . To show this, since we are allowed to take the measurements in any order, take the measurements in $M(L')$ before any other measurement. Then since the action of the unitary $U = U_p(\hat{\beta}, \hat{\gamma})$ on qubits in L' does not depend on any feature of the hypergraph outside of a radius of p around L' , the operators $M(L')$ are fully determined by the p -local neighborhoods of L' , and since we take them before every other measurement, the qubits are simply in their initial states when we make these measurements, thus the distribution of outputs is fully determined.

The same type of reasoning shows that the assignment to each $v \in V$ is statistically independent of the assignments to any set of vertices outside of a $2p$ -distance neighborhood of v . Take $L'' \subset V \setminus B(v, 2p)$. Then $M(v)$ acts on a radius- p ball around v , and each measurement in $M(L'')$ acts on a radius- p ball around a vertex in L'' , and by taking $\{M(v)\} \cup M(L'')$ before any of the measurements in $M(\{B(v, 2p) \setminus \{v\}\})$, we ensure that the qubits being measured by $M(v)$ are disjoint from and unentangled with those measured by anything in $M(L'')$. Hence the measurement $M(v)$ is independent of all measurements in $M(L'')$. We conclude that QAOA_p is a generic p -local algorithm. $\square$

3.3 Shared randomness between runs of a generic local algorithm

We describe a process to sample the outputs of a generic local algorithm when run twice on two different hypergraphs, so that the two runs of the algorithm can share randomness when the hypergraphs have some hyperedges in common.

This is not meant as a constructive algorithm, but a statistical process with no guarantee of feasible implementation.

The idea is to start with two t -coupled hypergraphs, which for large enough n , are likely to have some set of vertices L^+ whose p -neighborhoods are identical between the two hypergraphs. Since these vertices have identical p -neighborhoods, a generic p -local algorithm behaves identically on the vertices in L^+ . We pick a random t^+ fraction of the elements of L^+ , and assign the same labels to those vertices in the two coupled instances. Then the remaining labels on each hypergraph are assigned by generic p -local algorithms, conditioned on the output being consistent with the already assigned labels.

Definition 3.4 (Randomness-sharing for generic local algorithms).

Let A be a generic p -local algorithm, $(G_1, G_2) \sim \mathcal{H}_{d,k,n,t}$, and S be a label set. A pair of runs with t^+ -shared randomness of A on $G_1 = (V, E \cup E_1)$ and $G_2 = (V, E \cup E_2)$ with some shared edge set E is defined follows:

1. Let L^+ be the set of all vertices $v \in V$ such that $E(B_{G_1}(v, p)) \subset E$ and $E(B_{G_2}(v, p)) \subset E$. Generate the vertex set $L \subset L^+$ by including each element of L^+ independently with probability t^+ .
2. Since $\bigcup_{v \in L} (B_{G_1}(v, p)) = \bigcup_{v \in L} (B_{G_2}(v, p))$, the algorithm has the same joint marginal distribution for its outputs on L when it is run on G_1 or G_2 . Let $\sigma \in S^L$ be a sample from this joint marginal distribution.
3. Let σ_1 be a sample of $A(G_1)$, conditioned on $(\sigma_1)_v = \sigma_v$ for all $v \in L$. Similarly for σ_2 being a conditioned sample of $A(G_2)$. Then σ_1 and σ_2 are individually distributed the same as independent runs of the algorithm on G_1 and G_2 respectively, and together are the output of the two runs with t^+ -shared randomness.

4 Main Theorems

We formally state our main theorems and give the informal proof sketches in this section. The formal proofs are provided in [Section 6](#). First, let us specify the choice of parameters we are going to work with in the rest of the paper.

Parameter 4.1.

For every even $k \geq 4$, there exists $\eta_0 > 0$ such that the following holds: for every $\tau \in (0, 1)$ and $0 < \eta \leq \eta_0$, there exist $d_0, n_0 > 0$ from [Theorem 2.19](#) and we consider running a generic p -local algorithm on a random d -sparse k -uniform hypergraph G with size $n \geq n_0$, degree $d \geq d_0$ and p satisfying,

$$2p + 1 \leq \frac{(1 - \tau) \log n}{\log\left(\frac{d(k-1)}{\ln 2}\right)}.$$

4.1 Obstruction for generic local algorithms on diluted k -spin glasses

Theorem 4.2 (Obstruction theorem for diluted k -spin glasses).

Let (k, η, τ, d, n, p) be parameters satisfying [Parameter 4.1](#) and $\gamma > 0$ from [Theorem 5.1](#). Then, on running a generic p -local algorithm on a random d -sparse k -uniform hypergraph G with size n and degree d , the probability that the algorithm will output an assignment that is at least $(1 - \eta)$ -optimal is no more than $e^{-O(n^\gamma)}$.

Sketch of Proof: The proof follows the coupled interpolation argument in [[CGP⁺19](#), Section 3.3], and we sketch it briefly - The expected overlap between coupled solutions is continuous ([Lemma 6.1](#)) with the overlap being less than a at $t = 0$ with high probability if the solutions are nearly optimal ([Lemma 6.3](#)). The overlap is 1 at $t = 1$ ([Lemma 6.4](#)). Concentration of the overlap for any value of t is then shown by invoking [Theorem 5.1](#). The intermediate value theorem then immediately yields a contradiction to the *coupled* Overlap Gap Property ([Theorem 2.19](#)). This completes the proof for the obstruction. $\square$

4.2 Obstructions for generic local algorithms on (k, d) -CSP(f) with coupled OGP

Theorem 4.3 (Obstructions for (k, d) -CSP(f) with coupled OGP).

Let (k, η, τ, d, n, p) be parameters satisfying [Parameter 4.1](#) and $\gamma > 0$ from [Theorem 5.1](#). Let Ψ be a random problem instance of a signed or unsigned (k, d) -CSP(f) constructed as in [Definition 1.4](#) that satisfies a coupled OGP — That is, the hypergraph encoding G_Ψ of Ψ satisfies [Theorem 2.19](#) with the only difference that k can be any number ≥ 2 — as well as having the property that the $(1 - \eta_0)$ -multiplicatively optimal pairs of solutions to two independent instances of the CSP have overlap no more than the lower bound (a) of the OGP. Let σ be the output of a generic p -local algorithm on G_Ψ . Then the following holds:

$$\Pr_{\Psi, \sigma}[\text{val}_\Psi(\sigma) \geq (1 - \eta_0) \text{val}_\Psi] \leq e^{-O(n^\gamma)}.$$

Sketch of Proof: Once again, we first encode the problem instance Ψ into a representative hypergraph G_Ψ . Note that, by definition, the encoded instance G_Ψ satisfies a coupled OGP as

stated in [Theorem 2.19](#) over the underlying coupled interpolation stated in [Definition 2.18](#). The concentration of the hamming weight $|\sigma|$ of the solution is established by [Corollary 5.10](#) and the concentration of the objective value val_Ψ is established by [Theorem 4.5](#). The concentration of overlap of solutions for coupled instances over the interpolation specified in [Definition 2.18](#) is established by [Theorem 5.1](#). Then, by an argument similar to the one in the proof sketch of [Theorem 4.2](#), the obstruction follows. $\square$

Corollary 4.4 (Obstructions for random Max- k -XOR).

Let (k, η, τ, d, n, p) be parameters satisfying [Parameter 4.1](#) and $\gamma > 0$ from [Theorem 5.1](#). Let Ψ be a random problem instance of Max- k -XOR with $k \geq 4$ even. Let σ be the output of a generic p -local algorithm on G_Ψ . Then the following holds:

$$\Pr_{\Psi, \sigma}[\text{val}_\Psi(\sigma) \geq (1 - \eta_0)\text{val}_\Psi] \leq e^{-O(n^\gamma)}.$$

Proof. Combine [Theorem 4.3](#) with [Theorem 8.12](#) and [Lemma 8.14](#). $\square$

4.3 Concentration of objective function values of QAOA

For $p = O(1)$, concentration of the objective function value output by $\text{QAOA}_p(\hat{\beta}, \hat{\gamma})$ for sparse random constraint satisfaction problems (CSPs) is shown in [\[BBF⁺18\]](#). We state a result below which extends this to QAOA_p at depth $p < g(d, k) \log n$. While we state the result for $\text{QAOA}_p(\hat{\beta}, \hat{\gamma})$ specifically, this result will apply to any generic local algorithm. [\[BBF⁺18\]](#) cite a barrier in applying their techniques to QAOA_p at depth greater than $p = O(1)$ due to the limitation of McDiarmid's inequality as stated. We overcome this limitation by strengthening the inequality ([Lemma 7.6](#)) for highly biased distributions. This confirms the prediction of [\[BBF⁺18\]](#) about the "landscape independence" of $\text{QAOA}_p(\hat{\beta}, \hat{\gamma})$ at depth greater than $p = O(1)$.

Theorem 4.5 (QAOA_p landscape independence at $p < g(d, k) \log n$).

Let (k, η, τ, d, n, p) be parameters satisfying [Parameter 4.1](#). Let $U(\hat{\beta}, \hat{\gamma})$ be the unitary for a $\text{QAOA}_p(\hat{\beta}, \hat{\gamma})$ circuit. Furthermore, let the hamiltonian H_Ψ encode a problem instance Ψ of a (k, d) -CSP(f) constructed as in [Definition 1.4](#), such that,

$$H_\Psi = \sum_{i=1}^{|E|} H_i,$$

where each H_i is a k -local hamiltonian encoding f for the i -th clause. Then, the output $|\psi\rangle = U(\hat{\beta}, \hat{\gamma})|s\rangle$, where $|s\rangle$ is a symmetric product state, has an objective value that concentrates around the expected value as,

$$\Pr_{\Psi, |\psi\rangle} [|\langle \psi | H_\Psi | \psi \rangle - \mathbb{E}_{\Psi, |\psi\rangle} [\langle \psi | H_\Psi | \psi \rangle]| \geq \epsilon \cdot n] \leq e^{-O(n^\gamma)}, \forall \epsilon > 0.$$

Sketch of Proof: We encode the problem instance Ψ into a hypergraph G_Ψ . The constraint function f for every clause is set to be the local energy functions h_i on the appropriate k -subset of variables. Then [Theorem 5.3](#) and [Theorem 5.4](#) show concentration. $\square$

5 Concentration Analysis for Generic Local Algorithms

The most technical part of this work is to establish concentration theorems for generic local algorithms. Recall from [Figure 1](#) that to get obstruction from the coupled OGP, we have to show that the correlation between the outputs of a generic local algorithm on t -coupled hypergraphs is highly concentrated around its expected value $R(t)$. This is formally stated in the following theorem.

Theorem 5.1 (Generic local algorithm's outputs overlap on coupled hypergraphs).

Let (k, η, τ, d, n, p) be parameters satisfying [Parameter 4.1](#). When two random t -coupled hypergraphs $(G_1, G_2) \sim \mathcal{H}_{d,k,n,t}$ are sampled and a pair of t -shared-randomness runs of a generic p -local algorithm are made on G_1 and G_2 , the overlap between the respective outputs σ_1 and σ_2 concentrates. That is, $\forall \delta' > 0$, $\exists \gamma' > 0$, such that,

$$\Pr_{G_1, G_2, \sigma_1, \sigma_2} \left[\left| \langle \sigma_1, \sigma_2 \rangle - \mathbb{E}_{G_1, G_2, \sigma_1, \sigma_2} [\langle \sigma_1, \sigma_2 \rangle] \right| \geq \delta' \cdot n \right] \leq 2e^{-\delta' n \gamma'},$$

where the probability and expectation are over both the random sample of hypergraphs and the randomness of the algorithm.

To prove [Theorem 5.1](#), we have to show concentration with respect to both the internal randomness of the algorithm and the randomness from the problem instances. It turns out that the former is quite non-trivial due to the correlation between coupled hypergraphs as well as the dependencies introduced by each round of the local algorithm. This results in a generalization and strengthening of [\[FGG20a, Concentration Theorem\]](#) and [\[CGP⁺19, Lemmas 3.1 & 3.2\]](#).

To resolve the correlation issue, we introduce the notion of *locally mixed random vectors* ([Definition 5.2](#)) that capture the shared randomness between different runs of the local algorithm. We then show in [Theorem 5.3](#) that the correlation between a locally mixed random vector and the output of a generic p -local algorithm will still concentrate around its expectation with high probability.

Definition 5.2 (Locally mixed random vectors).

For a hypergraph G with vertex set $[n]$, a vector $r \in \mathbb{R}^m$ is a (G, p) -mixed random vector over $S_1, \dots, S_m \subseteq [n]$ with respect to $\sigma \in \mathbb{R}^n$ if:

- $|r_j| \leq 1$ for all $j \in [m]$,
- r_j is jointly independent of $r_{j'}$ for j' where $\text{dist}(S_j, S_{j'}) > 2p$ as well as σ_i for i where $\text{dist}(\{i\}, S_j) > 2p$, where,

$$\text{dist}(U, V) = \min_{u \in U, v \in V} \text{dist}(u, v).$$

Remark The purpose of the r vector is to enable reasoning about functions of more than one run of the algorithm, possibly with shared randomness between the runs. When considering functions of the output of a single run of the algorithm, it will suffice to take $r_j \equiv 1$, which is trivially (G, p) -mixed over $S_1, \dots, S_m$ with respect to σ for all G, p, σ , and $S_1, \dots, S_m$.

Theorem 5.3 (Concentration of local functions of spin configurations).

Let (k, η, τ, d, n, p) be parameters satisfying [Parameter 4.1](#). Let $\ell, m \in \mathbb{N}$. Let σ be the output of a generic p -local algorithm on a fixed hypergraph G . Let $v_{j,i} \in [n]$ for $j \in [m]$ and $i \in [\ell]$. Let $r = (r_1, \dots, r_m)$ be a (G', p) -mixed random vector ([Definition 5.2](#)) over $\{v_{1,i} \mid i \in [\ell]\}, \dots, \{v_{m,i} \mid i \in [\ell]\}$ with respect to σ for some hypergraph G' for which G is a subgraph of G' . Now, consider a sum,

$$X = \sum_{j \in [m]} h(\sigma_{v_{j,1}}, \dots, \sigma_{v_{j,\ell}}) r_j,$$

where $|h| \leq 1$. Suppose that each vertex v occurs at most C times among the different $v_{j,i}$. Then, provided that $B_{G'}(i, 2p)$ has at most n^A vertices in it for each $i \in [n]$, the following holds:

$$\Pr_{\sigma, r}[|X - \mathbb{E}[X]| \geq \delta n] \leq e^{-\Omega(\delta^2 m / (C \ell n^A))}.$$

Next, we show concentration (over the randomness of the problem instances) of functions of hypergraphs which satisfy a bounded-differences inequality with respect to small changes in the hypergraphs. This lemma is itself an application of the strengthening of McDiarmid's inequality, stated in [Lemma 7.6](#).

Theorem 5.4 (Concentration of bounded local differences on coupled hypergraphs).

Let (k, η, τ, d, n, p) be parameters satisfying [Parameter 4.1](#). Let A and a be the corresponding exponents from [Lemma 2.20](#). Let f be a function of two hypergraphs over n vertices V , such that

$$|f(G_1, G_2) - f(G'_1, G'_2)| \leq r(n)$$

for some r whenever (G_1, G_2) differs from (G'_1, G'_2) by the addition and/or removal of a single hyperedge $e \in [n]^k$ from one or both graph and

$$\max_{i \in [n]} \max_{G \in \{G_1, G_2, G'_1, G'_2\}} |B_G(i, p)| \leq n^{A/2}.$$

Then

$$\Pr_{G_1, G_2 \sim \mathcal{H}_{n,k,d,t}} \left[\left| f(G_1, G_2) - \mathbb{E}_{G_1, G_2} f(G_1, G_2) \right| \geq \delta n r(n) \right] \leq 2 \exp \left(\frac{-\delta^2 n}{4(2-t)d/k + 2\delta/3} \right) + 2 \exp(-n^{a/2}).$$

The theorem above is a generalization of the second part of [\[FGG20a, Concentration Theorem\]](#).

Organization of this section. In the rest of this section, we prove [Theorem 5.3](#) and [Theorem 5.4](#) in [Section 5.1](#) and [Section 5.2](#) respectively. Finally, we present the proof for [Theorem 5.1](#) in [Section 5.3](#) and show several useful corollaries in [Section 5.4](#).

5.1 Concentration over the internal randomness of the algorithm

In this subsection, we prove [Theorem 5.3](#) (restated below) which shows the concentration over the internal randomness of the algorithm. The proof is based on a Chernoff-style argument with a careful analysis on the combinatorial structure of the moment generating function.

Theorem 5.3 (Concentration of local functions of spin configurations).

Let (k, η, τ, d, n, p) be parameters satisfying [Parameter 4.1](#). Let $\ell, m \in \mathbb{N}$. Let σ be the output of a generic p -local algorithm on a fixed hypergraph G . Let $v_{j,i} \in [n]$ for $j \in [m]$ and $i \in [\ell]$. Let $r = (r_1, \dots, r_m)$ be a (G', p) -mixed random vector ([Definition 5.2](#)) over $\{v_{1,i} \mid i \in [\ell]\}, \dots, \{v_{m,i} \mid i \in [\ell]\}$ with respect to σ for some hypergraph G' for which G is a subgraph of G' . Now, consider a sum,

$$X = \sum_{j \in [m]} h(\sigma_{v_{j,1}}, \dots, \sigma_{v_{j,\ell}}) r_j,$$

where $|h| \leq 1$. Suppose that each vertex v occurs at most C times among the different $v_{j,i}$. Then, provided that $B_{G'}(i, 2p)$ has at most n^A vertices in it for each $i \in [n]$, the following holds:

$$\Pr_{\sigma, r}[|X - \mathbb{E}_{\sigma, r}[X]| \geq \delta n] \leq e^{-\Omega(\delta^2 m / (C \ell n^A))}.$$

Proof. We begin the proof by centering the variables X_j so that we can crucially conclude that they contribute 0 to the moment generating function when their expected value is taken. For technical reasons (to make odd moments zero), we also introduce a global independent random sign $s \sim \{\pm 1\}$.

$$Z_j = s \left[X_j - \mathbb{E}_{\sigma, r}[X_j] \right] = s \left[h(\sigma_{v_{j,1}}, \dots, \sigma_{v_{j,\ell}}) r_j - \mathbb{E}_{\sigma, r}[h(\sigma_{v_{j,1}}, \dots, \sigma_{v_{j,\ell}}) r_j] \right], \forall j \in [m].$$

Note as an immediate consequence that $\mathbb{E}_{\sigma, r}[Z_j] = 0, \forall j \in [m]$. Also, the goal now becomes showing $\Pr_{\sigma}[\sum_j Z_j \geq \delta m] < e^{-\Omega(\delta^2 m / (C \ell n^A))}$.

We start with analyzing the moment generating function of Z_j as follows,

$$\mathbb{E}_{\sigma, r} \left[\left(\sum_{j \in [m]} Z_j \right)^t \right] = \sum_{1 \leq j_1, \dots, j_t \leq m} \mathbb{E}_{\sigma, r} \left[\prod_{k \in \{j_1, \dots, j_t\}} Z_k \right], \forall t \in \mathbb{N}.$$

For the j -th summand on the right to be non-zero, every factor Z_k must be statistically dependent on at least one other term in $\{Z_{j_1}, \dots, Z_{j_t}\}$.

We count the number of summands that can be non-zero. If any Z_{j_i} is independent from all other $Z_{j_{i'}}$, then the entire term is zero, since $\mathbb{E}[Z_{j_i}] = 0$. Therefore if a summand is non-zero, then the interference graph between the factors Z_{j_i} has no isolated vertices. As a relaxation of this condition, the interference graph contains a forest with at most $t/2$ components as a subgraph. Therefore, we can upper bound the number of non-zero terms by summing over all such forests the number of ways to assign factors Z_{j_i} into vertices of that forest so that any two connected factors are not independent from each other.

We will split up the forests by the number of components they have, so first we count the number of forests over t vertices with w components for $w \leq t/2$. By a generalization of Cayley's formula [\[Tak90\]](#), there are $w t^{t-w-1}$ such forests if we assume that the first w vertices are in different components. Since we do not have the corresponding requirement on our factors, we may simply

choose w vertices arbitrarily to be in different components, multiplying by $\binom{t}{w}$ to find that there are at most $\binom{t}{w} w t^{t-w-1} m^w (C \ell n^A)^{t-w}$ ways to draw a forest with w components over our t indices $j_1, \dots, j_t$.

We now count the number of ways to assign the m factors into these forests so that any two factors connected by an edge are statistically dependent on each other. A vertex may be arbitrarily selected from each component of the factor graph (say, choose the one with the lowest index in $[t]$) to be assigned any of the m factors. Once that vertex has been assigned, each of its neighboring vertices has at most $C \ell n^A$ choices of dependent factors, since each factor is dependent on at most ℓn^A of the spins σ_i , and there are at most C factors which are a function of each spin. The same applies to all other vertices in the component.

Therefore, there are at most

$$\sum_{w=1}^{t/2} \binom{t}{w} w t^{t-w-1} m^w (C \ell n^A)^{t-w} \quad (5.5)$$

ways to generate a possible non-zero term.

Let $a_w = \binom{t}{w} w t^{t-w-1} m^w (C \ell n^A)^{t-w}$, and we will find the index w^* that maximizes a_{w^*} by computing the ratio

$$\frac{a_w}{a_{w-1}} = \frac{\frac{t!}{w!(t-w)!} w m}{\frac{t!}{(w-1)!(t-w+1)!} (w-1) t C \ell n^A} = \frac{(t-w+1)m}{(w-1)t C \ell n^A}.$$

At this point, we do some casework. In the case where $t \leq m/(C \ell n^A)$, the above ratio is greater than 1 whenever $w < 1 + t/2$, so $w^* = t/2$, recalling that we only have even moments because Z_j contains a factor of a global sign $s \sim \{\pm 1\}$. In the case where $t \geq m/(C \ell n^A)$, the ratio is less than 1 whenever $w > ((1 + 1/t)m + C \ell n^A)/(m/t + C \ell n^A)$, which is implied by $w > 1 + m/(2C \ell n^A)$, so $w^* \leq 1 + m/(2C \ell n^A)$.

In the former case where $t \leq m/(C \ell n^A)$ and $w^* = t/2$, we have

$$a_{w^*} = \frac{1}{2} \binom{t}{t/2} t^{t/2} m^{t/2} (C \ell n^A)^{t/2} \leq 2^t t^{t/2} m^{t/2} (C \ell n^A)^{t/2}.$$

In the latter case where $t \geq m/(C \ell n^A)$ and $w^* > 1 + m/(2C \ell n^A)$,

$$\begin{aligned} a_{w^*} &= \binom{t}{w^*} w^* t^{t-w^*-1} m^{w^*} (C \ell n^A)^{t-w^*} \\ &\leq \frac{e^{w^*}}{(w^*)^{w^*}} w^* t^{t-1} m^{w^*} (C \ell n^A)^{t-w^*} \\ &\leq e^{w^*} t^{t-1} m (C \ell n^A)^{t-1} \\ &\leq e^{t/2} t^{t-1} m (C \ell n^A)^{t-1}. \end{aligned}$$

In either case,

$$a_{w^*} \leq 2^t t^{t/2} m^{t/2} (C \ell n^A)^{t/2} + e^{t/2} t^{t-1} m (C \ell n^A)^{t-1},$$

so the number of non-zero terms in the expression of $\mathbb{E}[(\sum Z_j)^t]$ is at most $t/2$ times this bound on the maximum value of a_{w^*} , and, recalling that $|Z_j| \leq 1$,

$$\mathbb{E}_\sigma \left[\left(\sum_{j \in [m]} Z_j \right)^t \right] \leq 2^{t-1} t^{t/2+1} m^{t/2} (C \ell n^A)^{t/2} + \frac{1}{2} e^{t/2} t^t m (C \ell n^A)^{t-1}.$$

We multiply both sides by θ^t and divide by $t!$ and sum over even t (recalling that the definition of Z_j contains a random global sign making all odd moments zero) to obtain a bound on the moment generating function

$$\mathbb{E}_\sigma \left[e^{\theta \sum_{j \in [m]} Z_j} \right] \leq \sum_{t \in \mathbb{N} \text{ even}} \frac{\theta^t}{t!} 2^{t-1} t^{t/2+1} m^{t/2} (C \ell n^A)^{t/2} + \sum_{t \in \mathbb{N} \text{ even}} \frac{\theta^t}{2 \cdot t!} e^{t/2} t^t m (C \ell n^A)^{t-1}.$$

We handle the two terms separately. For the first one, we reparameterize the index t and then make use of Stirling's approximation:

$$\begin{aligned} \sum_{t \in \mathbb{N} \text{ even}} \frac{\theta^t}{t!} 2^{t-1} t^{t/2+1} m^{t/2} (C \ell n^A)^{t/2} &= \sum_{t \in \mathbb{N}} \frac{\theta^{2t}}{(2t)!} 2^{2t-1} (2t)^{t+1} m^t (C \ell n^A)^t \\ &= \sum_{t \in \mathbb{N}} \frac{t^{t+1}}{(2t)!} (8\theta^2 m C \ell n^A)^t \\ &\leq \sum_{t \in \mathbb{N}} \frac{t^{t+1}}{\sqrt{2\pi} (2t)^{2t+1/2} e^{-2t}} (8\theta^2 m C \ell n^A)^t \\ &\leq \sum_{t \in \mathbb{N}} \frac{t^{t+1} t^{t+1/2}}{\sqrt{2\pi} (2t)^{2t+1/2} e^{-t-1} t!} (8\theta^2 m C \ell n^A)^t \\ &= \frac{e}{2\sqrt{\pi}} \sum_{t \in \mathbb{N}} \frac{t (2e\theta^2 m C \ell n^A)^t}{t!} \\ &= \frac{\theta^2 m C \ell n^A e^{2e\theta^2 m C \ell n^A + 2}}{\sqrt{\pi}}. \end{aligned}$$

For the second term,

$$\begin{aligned} \sum_{t \in \mathbb{N} \text{ even}} \frac{\theta^t}{2 \cdot t!} e^{t/2} t^t m (C \ell n^A)^{t-1} &\leq \sum_{t \in \mathbb{N}} \frac{\theta^t}{2 \cdot t!} e^{t/2} t^t m (C \ell n^A)^{t-1} \\ &\leq \frac{m}{2\sqrt{2\pi} C \ell n^A} \sum_{t \in \mathbb{N}} \frac{(e^{3/2} \theta C \ell n^A)^t}{\sqrt{t}} \\ &\leq \frac{m}{2\sqrt{2\pi} C \ell n^A} \sum_{t \in \mathbb{N}} (e^{3/2} \theta C \ell n^A)^t \\ &\leq \frac{m}{2\sqrt{2\pi} C \ell n^A (1 - e^{3/2} \theta C \ell n^A)}. \end{aligned}$$

Putting these bounds together,

$$\mathbb{E}_\sigma \left[e^{\theta \sum_{j \in [m]} Z_j} \right] \leq \frac{\theta^2 m C \ell n^A e^{2e\theta^2 m C \ell n^A + 2}}{\sqrt{\pi}} + \frac{m}{2\sqrt{2\pi} C \ell n^A (1 - e^{3/2} \theta C \ell n^A)}.$$

By Markov's inequality,

$$\Pr_{\sigma} \left[\sum Z_j \geq \delta m \right] \leq e^{-\theta \delta m} \mathbb{E}_{\sigma} \left[e^{\theta \sum_{j \in [m]} Z_j} \right]$$

for all θ . Choosing $\theta = \delta / (2e^{3/2} C \ell n^A)$ with $0 < \delta \leq 1$, we get

$$\mathbb{E}_{\sigma} \left[e^{\theta \sum_{j \in [m]} Z_j} \right] \leq \frac{\delta^2 m e^{\delta^2 m / (2e^2 C \ell n^A)}}{4e(C \ell n^A) \sqrt{\pi}} + \frac{m}{\sqrt{2\pi} C \ell n^A}.$$

Therefore,

$$\Pr_{\sigma} \left[\sum_j Z_j \geq \delta m \right] \leq \frac{\delta^2 m e^{(-\sqrt{e}+1)\delta^2 m / (2e^2 C \ell n^A)}}{4e(C \ell n^A) \sqrt{\pi}} + \frac{m e^{-\delta^2 m / (2e^{3/2} C \ell n^A)}}{\sqrt{2\pi} C \ell n^A}.$$

The identical bound holds for $\Pr_{\sigma} \left[\sum_j Z_j \leq -\delta m \right]$.

□

5.2 Concentration over randomness of problem instance

In this subsection, we prove [Theorem 5.4](#) which shows concentration of coupled hypergraphs with respect to the randomness of problem instances. We start with stating and proving a special case of [Theorem 5.4](#) to illustrate the structure of the argument in a simpler setting.

Lemma 5.6 (Concentration of Bounded Local Differences on Random Hypergraphs).

For every $p \in \mathbb{N}$, let A and a be the corresponding exponents from [Lemma 2.20](#). Let f be a function of a hypergraph on n vertices, such that $|f(G_1) - f(G_2)| \leq r(n)$ whenever $\max_{i \in [n]} \max(|B_{G_1}(i, p)|, |B_{G_2}(i, p)|) \leq n^{A/2}$ and G_1 differs from G_2 by the addition or removal of a single edge. Then

$$\Pr_{G \sim \mathcal{H}_{n,k,d}} \left[\left| f(G) - \mathbb{E}_G f(G) \right| \geq \delta n r(n) \right] \leq 2 \exp \left(\frac{-\delta^2 n}{4d/k + 2\delta/3} \right) + \exp \left(-n^{a/2} \right).$$

Proof. The proof follows by essentially the same arguments used in the second part of [\[FGG20a, Concentration Theorem\]](#), although we have to derive a strengthening of McDiarmid's inequality ([Lemma 7.6](#)), deferred to [Section 7](#).

Let K_n be the set of hypergraphs over n vertices V with small neighborhoods

$$K_n = \{G \in \mathcal{H}_{n,k} \mid \max_{v \in V} |B_G(v, p)| \leq n^{A/2}\}.$$

Let $\rho(G, G')$ be equal to $|E \triangle E'|$ where $\triangle$ is the symmetric difference, for $G = (V, E)$ and $G' = (V, E')$. Then let

$$g(G) = \min_{G' \in K_n} f(G') + \rho(G, G') r(n).$$

Now g has the property that $|g(G_1) - g(G_2)| \leq r(n)$ whenever G_1 differs from G_2 by the addition or removal of a single edge.

$G \sim \mathcal{H}_{n,k,d}$ may be viewed as the agglomeration of n^k different independent random variables, one for each possible hyperedge e , each random variable denoting the multiplicity of that hyperedge,

distributed as $\text{Poisson}(dn/n^k/k)$, since the sum of independent Poisson random variables is itself another Poisson random variable. And each of these variables is highly biased, being equal to 0 with probability $\exp(-dn/n^k/k) \geq 1 - dn/n^k/k$. Therefore, applying [Lemma 7.6](#) on g as a function of n^k independent variables, simplifying, and applying the bound $2 - dn/n^k/k \leq 2$,

$$\Pr_{G \sim \mathcal{H}_{n,k,d}} [|g(G) - \mathbb{E}_G g(G)| \geq \delta n r(n)] \leq 2 \exp\left(\frac{-\delta^2 n}{4d/k + 2\delta/3}\right).$$

Finally, $f(G) = g(G)$ whenever $G \in K_n$, and by [Lemma 2.20](#),

$$\Pr_{G \sim \mathcal{H}_{n,k,d}} [G \notin K_n] \leq e^{-n^{a/2}}.$$

Therefore, by a union bound,

$$\Pr_{G \sim \mathcal{H}_{n,k,d}} [|f(G) - \mathbb{E}_G f(G)| \geq \delta n r(n)] \leq 2 \exp\left(\frac{-\delta^2 n}{4d/k + 2\delta/3}\right) + \exp(-n^{a/2}).$$

□

We now prove [Theorem 5.4](#) (restated below) which shows that the expected values for certain functions on coupled hypergraphs also concentrate. More specifically, we will assert that all sufficiently *local* functions of pairs of hypergraphs will concentrate very heavily around the expected value of the function.

Theorem 5.4 (Concentration of bounded local differences on coupled hypergraphs).

Let (k, η, τ, d, n, p) be parameters satisfying [Parameter 4.1](#). Let A and a be the corresponding exponents from [Lemma 2.20](#). Let f be a function of two hypergraphs over n vertices V , such that

$$|f(G_1, G_2) - f(G'_1, G'_2)| \leq r(n)$$

for some r whenever (G_1, G_2) differs from (G'_1, G'_2) by the addition and/or removal of a single hyperedge $e \in [n]^k$ from one or both graph and

$$\max_{i \in [n]} \max_{G \in \{G_1, G_2, G'_1, G'_2\}} |B_G(i, p)| \leq n^{A/2}.$$

Then

$$\Pr_{G_1, G_2 \sim \mathcal{H}_{n,k,d,t}} \left[\left| f(G_1, G_2) - \mathbb{E}_{G_1, G_2} f(G_1, G_2) \right| \geq \delta n r(n) \right] \leq 2 \exp\left(\frac{-\delta^2 n}{4(2-t)d/k + 2\delta/3}\right) + 2 \exp(-n^{a/2}).$$

Proof. We mostly follow the proof of [Lemma 5.6](#).

Let $\rho(G_1, G_2, G'_1, G'_2) = |(E_1 \triangle E'_1) \cup (E_2 \triangle E'_2)|$, where $G_1 = (V, E_1)$, $G_2 = (V, E_2)$, $G'_1 = (V, E'_1)$, and $G'_2 = (V, E'_2)$. Then let

$$g(G_1, G_2) = \min_{G'_1 \in K_n, G'_2 \in K_n} f(G'_1, G'_2) + \rho(G_1, G_2, G'_1, G'_2) r(n),$$

so that $|g(G_1, G_2) - g(G'_1, G'_2)| \leq r(n)$ whenever (G_1, G_2) differs from (G'_1, G'_2) by the addition or removal of a single hyperedge in $[n]^k$ from one or both graphs.

To apply [Lemma 7.6](#), we consider g as a function of n^k variables, one for each possible hyperedge. The set of possible values for each variable is $\{\emptyset, \{G_1\}, \{G_2\}, \{G_1, G_2\}\}$, specifying which of the hypergraphs have that edge. In the coupled random hypergraph model according to [Definition 2.18](#), where here $G_1 = (V, E \cup E_1)$ and $G_2 = (V, E \cup E_2)$ each edge's multiplicity in E is given by a $\text{Poisson}(tdn/n^k/k)$ distribution, and its multiplicities in E_1 and E_2 are given by $\text{Poisson}((1-t)dn/n^k/k)$ distributions, for a total probability of $(1 - \exp(-tdn/n^k/k))(1 - \exp(-(1-t)dn/n^k/k))^2 \leq (2-t)dn/n^k/k$ that this edge is in either hypergraph. Therefore, applying [Lemma 7.6](#),

$$\Pr_{G_1, G_2 \sim \mathcal{H}_{n,k,d,t}} \left[\left| g(G_1, G_2) - \mathbb{E}_{G_1, G_2} g(G_1, G_2) \right| \geq \delta n r(n) \right] \leq 2 \exp \left(\frac{-\delta^2 n}{4(2-t)d/k + 2\delta/3} \right).$$

Finally, $f(G_1, G_2) = g(G_1, G_2)$ whenever $(G_1, G_2) \in K_n \times K_n$. Since the marginal distribution of G_1 in $G_1, G_2 \sim \mathcal{H}_{n,k,d,t}$ is the same as the distribution of $G_1 \sim \mathcal{H}_{n,k,d}$, and the same holds for G_2 , by a union bound,

$$\Pr_{G_1, G_2 \sim \mathcal{H}_{n,k,d,t}} [(G_1, G_2) \notin K_n \times K_n] \leq \Pr_{G_1 \sim \mathcal{H}_{n,k,d}} [G_1 \notin K_n] + \Pr_{G_2 \sim \mathcal{H}_{n,k,d}} [G_2 \notin K_n] \leq 2e^{-n^{a/2}}.$$

Therefore, by another union bound,

$$\Pr_{G_1, G_2 \sim \mathcal{H}_{n,k,d,t}} \left[\left| f(G_1, G_2) - \mathbb{E}_{G_1, G_2} f(G_1, G_2) \right| \geq \delta n r(n) \right] \leq 2 \exp \left(\frac{-\delta^2 n}{4(2-t)d/k + 2\delta/3} \right) + 2 \exp(-n^{a/2}).$$

□

5.3 Proof of Theorem 5.1

Finally, we are ready to prove [Theorem 5.1](#) (restated below) using [Theorem 5.3](#) and [Theorem 5.4](#).

Theorem 5.1 (Generic local algorithm's outputs overlap on coupled hypergraphs).

Let (k, η, τ, d, n, p) be parameters satisfying [Parameter 4.1](#). When two random t -coupled hypergraphs $(G_1, G_2) \sim \mathcal{H}_{d,k,n,t}$ are sampled and a pair of t -shared-randomness runs of a generic p -local algorithm are made on G_1 and G_2 , the overlap between the respective outputs σ_1 and σ_2 concentrates. That is, $\forall \delta' > 0$, $\exists \gamma' > 0$, such that,

$$\Pr_{G_1, G_2, \sigma_1, \sigma_2} \left[\left| \langle \sigma_1, \sigma_2 \rangle - \mathbb{E}_{G_1, G_2, \sigma_1, \sigma_2} [\langle \sigma_1, \sigma_2 \rangle] \right| \geq \delta' \cdot n \right] \leq 2e^{-\delta' n^{\gamma'}},$$

where the probability and expectation are over both the random sample of hypergraphs and the randomness of the algorithm.

Proof. By [Theorem 5.3](#), taking $G = G_2$, $m = n$, $C = \ell = 1$, $v_{j,1} = j$, $h(s) = s$, and $r_j = u_j$,

$$\Pr_{\sigma_2} \left[\left| \langle \sigma_2, u \rangle - \mathbb{E}_{\sigma_2} \langle \sigma_2, u \rangle \right| \geq 2\delta n \right] \leq \exp\left(-\Omega(\delta^2 n^{1-A})\right),$$

for every bounded vector $u \in [-1, 1]^n$, with the probability and expectation over the randomness of the algorithm, where A corresponds to the exponent in [Lemma 2.20](#).

Again by [Theorem 5.3](#), taking $G = G_1$, $m = n$, $C = \ell = 1$, $v_{j,1} = j$, $h(s) = s$, $r_j = (\sigma_2)_j$, and $G' = G_1 \cup G_2$, using the fact that $\langle \sigma_1, \sigma_2 \rangle = \sum_{i \in [n]} (\sigma_1)_i (\sigma_2)_i$,

$$\Pr_{\sigma_1, \sigma_2} \left[\left| \langle \sigma_1, \sigma_2 \rangle - \mathbb{E}_{\sigma_1, \sigma_2} \langle \sigma_1, \sigma_2 \rangle \right| \geq 2\delta n \right] \leq \exp\left(-\Omega(\delta^2 n^{1-A})\right) + \exp\left(-\Omega(\delta^2 n^{1-A'})\right), \quad (5.7)$$

with the probability and expectation over the randomness of the algorithm. In the above, we utilize the fact that $G' \sim \mathcal{H}_{d(2-t), k, n}$. To bound A' , one bounds the $2p$ -neighborhood of G' using [Lemma 2.20](#) invoked with degree $d(2-t) \leq 2d$.

We argue now that $f(G_1, G_2) := \mathbb{E}_{\sigma_1, \sigma_2} \langle \sigma_1, \sigma_2 \rangle$ satisfies a bounded-differences inequality, so that $|f(G_1, G_2) - f(G'_1, G'_2)| \leq 2kn^A$ whenever $\max_{i \in [n]} \max_{G \in \{G_1, G_2, G'_1, G'_2\}} |B_G(i, p)| \leq n^A$ and $|(E(G_1) \triangle E(G'_1)) \cup (E(G_2) \triangle E(G'_2))| \leq 1$. This is because by the definition of p -local, the only coordinates of σ_1 or σ_2 that can change in marginal distribution when an edge is changed are those that are within a distance of p from any of the k endpoints of the changed hyperedge, in G_1 and G'_1 , or G_2 and G'_2 respectively.

Therefore by [Theorem 5.4](#),

$$\Pr_{G_1, G_2} \left[\left| \mathbb{E}_{\sigma_1, \sigma_2} \langle \sigma_1, \sigma_2 \rangle - \mathbb{E}_{G_1, G_2} \mathbb{E}_{\sigma_1, \sigma_2} \langle \sigma_1, \sigma_2 \rangle \right| \geq 2k\delta' n^A n \right] \leq 2 \exp\left(\frac{-(\delta')^2 n}{4(2-t)d/k + 2\delta'/3}\right) + 2 \exp\left(-n^{a/2}\right),$$

where a is from the statement of [Lemma 2.20](#).

By a union bound and triangle inequality with (5.7) then, and taking $\delta' = \delta/(kn^A)$,

$$\begin{aligned} \Pr_{G_1, G_2, \sigma_1, \sigma_2} \left[\left| \langle \sigma_1, \sigma_2 \rangle - \mathbb{E}_{G_1, G_2} \mathbb{E}_{\sigma_1, \sigma_2} \langle \sigma_1, \sigma_2 \rangle \right| \geq 4\delta n \right] &\leq \\ &2 \exp\left(\frac{-\delta^2 n}{4(2-t)kn^{2A}d + 2kn^A\delta/3}\right) + 2 \exp\left(-n^{a/2}\right) + \exp\left(-\Omega(\delta^2 n^{1-A})\right) + \\ &\exp\left(-\Omega(\delta^2 n^{1-A'})\right). \end{aligned}$$

□

5.4 Corollaries of concentration analysis

As a corollary to [Theorem 5.1](#), we can inherit the concentrated overlap of two *independent* hypergraphs $G_1, G_2 \sim \mathcal{H}_{d, k, n}$. This is essential to reasoning about the fact that the overlap at $t = 0$ between the solutions is desirably small with high probability.

Corollary 5.8 (Generic local algorithm output overlap on independent hypergraphs).

Suppose p is as in [Lemma 2.20](#). When two random hypergraphs G_1, G_2 are sampled i.i.d. from $\mathcal{H}_{d,k,n}$ and a generic p -local algorithm is run on both hypergraphs, the overlap between the respective solutions σ_1 and σ_2 concentrates. That is, $\exists \delta' > 0$ and $\gamma' > 0$, such that,

$$\Pr_{G_1, G_2, \sigma_1, \sigma_2} [|\langle \sigma_1, \sigma_2 \rangle - \mathbb{E}[\langle \sigma_1, \sigma_2 \rangle]| \geq \delta' \cdot n] \leq 2e^{-\delta' n^{\gamma'}}, \quad (5.9)$$

where the probability is over both the random sample of hypergraphs and the randomness of the algorithm.

Proof. By [Theorem 5.1](#) at $t = 0$ with G_1 independent of G_2 . $\square$

We now assert the concentration of hamming weight of the solutions output by a generic p -local algorithm over the randomness of the algorithm. This result is equivalent to the second part of [\[FGG20a, Concentration Theorem\]](#) and [\[CGP⁺19, Lemma 3.2\]](#). In our case, it is a corollary of [Theorem 5.3](#).

Corollary 5.10 (Concentration of hamming weight of generic local algorithm output).

Given a d -sparse k -uniform hypergraph $G \sim \mathcal{H}_{d,k,n}$, suppose that,

$$\Pr_G [\max_i B_G(v_i, 2p) \geq n^A] \leq e^{-n^a},$$

and

$$\Pr_G [\max_i B_G(v_i, p) \geq n^{\frac{A}{2}}] \leq e^{-n^{\frac{a}{2}}},$$

Let σ be the output of a generic p -local algorithm. Then, $\exists \gamma > 0$ such that, $\forall \delta > 0$,

$$\Pr_{G, \sigma} \left[\left| |\sigma| - \mathbb{E}_{G, \sigma} [|\sigma|] \right| \geq \delta n \right] \leq e^{-\delta^2 n^\gamma}, \quad (5.11)$$

where $|\sigma| = \sum_{i=1}^n \frac{1}{2}(1 - \sigma_i)$ is the number of -1 s in σ .

Proof. Instantiate [Theorem 5.3](#) by asserting that the following holds for the graph G ,

$$\Pr_\sigma \left[\left| \sum_{i=1}^n \sigma_i - \mathbb{E}_\sigma [\sigma] \right| \geq 2\delta n \right] \leq \exp(-\Omega(\delta^2 n^{1-A})). \quad (5.12)$$

with $\ell = 1$, $C = 1$, $m = n$, $v_{j,1} = j$, $r_j = 1$, and $h(s) = s$. We now invoke [Lemma 5.6](#) with $f = \mathbb{E}_\sigma [|\sigma|]$. Now, notice that the definition of generic p -local algorithms ([Definition 3.1](#)) implies that on adding or removing a hyperedge e , f can change by no more than $2kn^A$, since changes are restricted to the p -neighborhood. Therefore, provided G is a graph with $\max_i B(v_i, 2p) \leq n^A$, we obtain,

$$\Pr_{G, \sigma} \left[\left| \mathbb{E}_\sigma [|\sigma|] - \mathbb{E}_G \mathbb{E}_\sigma [|\sigma|] \right| \geq 2k\delta' n^A n \right] \leq 2 \exp \left(\frac{-\delta'^2 n}{4d/k + 2\delta'/3} \right) + \exp(-n^{a/2}).$$

By setting $\delta' = \delta/kn^A$, and taking a triangle inequality over Equation 5.12 and a union bound,

$$\Pr_{G,\sigma} \left[\left| |\sigma| - \mathbb{E}_{G,\sigma} [|\sigma|] \right| \geq 2\delta n \right] \leq 2 \exp \left(\frac{-\delta^2 n^{1-2A}}{4dk + 2\delta k/(3n^A)} \right) + \exp(-n^{a/2}) + \exp(-\Omega(\delta^2 n^{1-A})).$$

□

The last corollary we obtain from Theorem 5.3 is that the objective function value, that is, the energy corresponding to the spin configuration output by a generic p -local algorithm, also concentrates heavily around the expected value. This is equivalent to [CGP⁺19, Lemma 3.1] with substantially stronger concentration.

Corollary 5.13 (Concentration of energy of generic local algorithm output).

Given a d -sparse k -uniform hypergraph $G \sim \mathcal{H}_{d,k,n}$, suppose that

$$\Pr_G [\max_i B_G(v_i, 2p) \geq n^A] \leq e^{-n^a},$$

and

$$\Pr_G [\max_i B_G(v_i, p) \geq n^{\frac{A}{2}}] \leq e^{-n^{\frac{a}{2}}}.$$

Let σ be the output of a generic p -local algorithm. Then, $\exists \gamma > 0$, such that, $\forall \delta > 0$,

$$\Pr_{G,\sigma} \left[\left| H_{d,k,n}^G(\sigma) - \mathbb{E}_{G,\sigma} [H_{d,k,n}^G(\sigma)] \right| \geq \delta n \right] \leq e^{-\delta^2 n^\gamma}.$$

Proof. For sufficiently large n , the degree of the vertices in G are distributed as $\text{Poisson}(d)$. Therefore, by a standard Chernoff bound for the Poisson distribution [Can17, Theorem 1],

$$\Pr_{\text{Poisson}(d)} \left[\deg(v) \geq d + n^{\frac{a}{2}} \right] \leq \exp \left(\frac{-n^a}{2(d + n^{a/2})} \right).$$

Applying a union bound to the above for every vertex $v \in V(G)$ yields that the degree of every vertex can be upper bounded by $d + n^{\frac{a}{2}} \leq 2n^{\frac{a}{2}}$, for sufficiently large n .

An instantiation of Theorem 5.3 on a graph G chosen so that the degree of every vertex is not more than $2n^{\frac{a}{2}}$ implies that,

$$\Pr_\sigma \left[\left| H_{d,k,n}^G(\sigma) - \mathbb{E}_\sigma [H_{d,k,n}^G(\sigma)] \right| \geq 2\delta n \right] \leq \exp \left(-\Omega \left(\frac{\delta^2 d n^{1-A-a/2}}{2k} \right) \right), \quad (5.14)$$

where $m = dn$, $\ell = k$, $C = 2n^{\frac{a}{2}}$ and $h_j(s_{j,1}, \dots, s_{j,k}) = -\prod_{i=1}^\ell s_{j,i}$, where $s_{j,i}$ denotes the spin of the i -th vertex of the j -th hyperedge.

We instantiate Lemma 5.6 with $f = \mathbb{E}_\sigma [H_{d,k,n}]$. Note that if we change one hyperedge e of G , then because of the fact that a generic p -local algorithm will only affect the joint distribution in a p -neighborhood around every vertex, we notice that f cannot change by more than $2Ckn^A = 4kn^{A+\frac{a}{2}}$

when G satisfies the bounded neighborhood proposition in the hypothesis. Therefore,

$$\Pr_{G,\sigma} \left[\left| \mathbb{E}_{\sigma} [H_{d,k,n}^G(\sigma)] - \mathbb{E}_G \mathbb{E}_{\sigma} [H_{d,k,n}^G] \right| \geq 2Ck\delta' n^A n \right] \leq 2 \exp \left(\frac{-\delta'^2 n}{4d/k + 2\delta'/3} \right) + \exp \left(-n^{a/2} \right) + \exp \left(\frac{-n^a}{2(d + n^{a/2})} \right).$$

By a union bound and triangle inequality with [Equation 5.14](#), followed by taking $\delta' = \delta/(Ckn^A)$,

$$\Pr_{G,\sigma} \left[\left| H_{d,k,n}^G(\sigma) - \mathbb{E}_{G,\sigma} [H_{d,k,n}^G(\sigma)] \right| \geq 4\delta n \right] \leq 2 \exp \left(\frac{-\delta^2 n^{1-2A-a}}{4dk + \frac{2}{3}k\delta/n^{A+a/2}} \right) + \exp \left(-n^{a/2} \right) + \exp \left(\frac{-n^a + \log(n)}{2(d + n^{a/2})} \right) + \exp \left(\frac{-\Omega(\delta^2 d n^{1-A-a/2})}{2k} \right).$$

Note that we can guarantee that $2a + A < 1$ with an appropriate choice of $a, A > 0$, which is made possible by choosing an appropriate value of $\tau > 0$ in [Parameter 4.1](#). For an explicit characterization, refer to [Appendix B](#) and [\[FGG20a, Equation \(75\), Neighborhood Size Theorem\]](#). $\square$

6 Proofs of main theorems

We conclude by putting the previous results together to prove our main results. As stated prior, we work in the setting of [Parameter 4.1](#).

6.1 Proof of Theorem 4.2

We establish that, as a function of t , the expected overlap between solutions output by runs of a generic local algorithm on t -coupled instances G_1 and G_2 with t -shared randomness is continuous.

Lemma 6.1 (Continuity of expected overlap).

If $G_1, G_2 \sim \mathcal{H}_{d,k,n,t}$, and σ_1, σ_2 are the random outputs of t -shared randomness runs of a generic p -local algorithm on each hypergraph,

$$\mathbb{E}_{G_1, G_2, \sigma_1, \sigma_2} [R_{n,t}(\sigma_1, \sigma_2)] \text{ is a continuous function of } t \text{ for all } t \in (0, 1).$$

Proof. We express using linearity of expectation

$$\mathbb{E}_{G_1, G_2, \sigma_1, \sigma_2} [R_{n,t}(\sigma_1, \sigma_2)] = \frac{1}{n} \sum_{i \in [n]} \mathbb{E}_{G_1, G_2, \sigma_1, \sigma_2} [(\sigma_1)_i (\sigma_2)_i] = \mathbb{E}_{G_1, G_2, \sigma_1, \sigma_2} [(\sigma_1)_1 (\sigma_2)_1].$$

Let $s(H_1, H_2)$ be the expectation value $\mathbb{E}_{\sigma_1, \sigma_2} (\sigma_1)_1 (\sigma_2)_1$ when σ_1 and σ_2 are the results of a pair of runs of the p -local algorithm with t -shared randomness on H_1 and H_2 respectively. Then

$$\mathbb{E}_{G_1, G_2, \sigma_1, \sigma_2} [R_{n,t}(\sigma_1, \sigma_2)] = \sum_{H_1, H_2} \Pr_{G_1, G_2 \sim \mathcal{H}_{n,k,d,t}} [H_1 = G_1 \wedge H_2 = G_2] s(H_1, H_2).$$

Since the probability of each combination of G_1 and G_2 being sampled is a continuous function of t , so then is $\mathbb{E}_{G_1, G_2, \sigma_1, \sigma_2} [R_{n,t}(\sigma_1, \sigma_2)]$. $\square$

Note that at $t = 0$, the hypergraphs $G_1, G_2 \sim \mathcal{H}_{d,k,n,t}$ are independent of each other. We will show that there is very little overlap between the outputs σ_1, σ_2 of a generic local algorithm on each instance, if both solutions σ_1 and σ_2 are better approximations to the optimum value by a factor of $(1 - \frac{\eta_0}{P(k)})$. This follows directly from a combination of [CGP⁺19, Lemma 3.3] and Theorem 5.3. We begin by restating [CGP⁺19, Lemma 3.3].

Lemma 6.2 (Hamming weight of near-optimal solutions, [CGP⁺19, Lemma 3.3]).

Given two independent and uniformly random hypergraphs $G_1, G_2 \sim \mathcal{H}_{d,k,n}$, and solutions σ_1 and σ_2 , such that,

$$\frac{H^{G_i}(\sigma_i)}{n} > M(k, d) \left(1 - \frac{\eta_0}{P(k)}\right),$$

for $i \in \{1, 2\}$ and any $\eta_0 > 0$. Then, $|\sigma_i| < d^{-\frac{1}{2k}} (4\eta_0)^{\frac{1}{k}}$ with probability $\geq 1 - O(e^{-n})$.

The above lemma allows us to bound the overlap between two instances when presented in conjunction with a concentration argument about the hamming weight of instances produced by a generic local algorithm.

Lemma 6.3 (Small overlap at $t = 0$).

For any two hypergraphs $G_1, G_2 \sim \mathcal{H}_{d,k,n,0}$, let σ_1, σ_2 be the random outputs of two 0-shared randomness runs of a generic p -local algorithm on the instances. If σ_1 and σ_2 satisfy the optimality of Lemma 6.2, then

$$R_{n,0}(\sigma_1, \sigma_2) \leq d^{-\frac{1}{k}},$$

with probability $\geq 1 - e^{-O(n^\gamma)}$.

Proof. The proof is essentially the same as that in [CGP⁺19, Section 3.3] with the exception that we provide stronger concentration via Theorem 5.3.

By Lemma 6.2, we have that,

$$\left| \frac{1}{n} \sum_{i=1}^n \sigma_{j_i} \right| < d^{-1/2k} (4\eta_0)^{1/k}, \text{ for } j = 1, 2,$$

with probability $\geq 1 - O(e^{-n})$. Choose $\eta_0 = \frac{1}{8}$, per [CGP⁺19]. This implies,

$$|\sigma_i| \leq d^{-1/2k} 2^{-1/k} \leq d^{-1/2k}.$$

Furthermore, because the output of generic p -local algorithms concentrate,

$$\mathbb{E}_{G_i, \sigma_i} [|\sigma_i|] \leq d^{-1/2k},$$

with probability $\geq 1 - e^{-O(n^\gamma)}$ (as implied by Corollary 5.10). Now, note that at $t = 0$, the generic p -local algorithm is equally likely to generate $\sigma_i, i \in \{1, 2\}$ with a given hamming-weight (in the ± 1 basis). Consequently, by independence, the largest expected overlap is the square of the maximum possible hamming weight. This yields,

$$R_{n,0}(\sigma_1, \sigma_2) \leq (d^{-1/2k})^2 = d^{-1/k}.$$

To establish that this is less than a , simply choose $d \geq \max(d_0, a^{-k})$. □

We now assert that the output of two generic local algorithm runs with fully shared randomness on fully coupled (and therefore identical) hypergraphs G_1 and G_2 has overlap 1.

Lemma 6.4 (Large overlap at $t = 1$).

For any two hypergraphs $G_1, G_2 \sim \mathcal{H}_{d,k,n,1}$, given that σ_1, σ_2 are the outputs of two 1-shared randomness runs of a p -local algorithm on the instances,

$$R_{n,1}(\sigma_1, \sigma_2) = 1.$$

Proof. By [Definition 2.18](#), at $t = 1$, the two hypergraphs G_1 and G_2 sampled from $\mathcal{H}_{d,k,n,1}$ are in fact the same. Therefore every vertex's neighborhood is the same in G_1 and G_2 , and so with $t = 1$ in [Definition 3.4](#), $L^+ = V$ and so $\sigma = \sigma_1 = \sigma_2$ and the two outputs are identical. $\square$

The arguments above immediately yield a contradiction to the coupled OGP stated in [Theorem 2.19](#) via the intermediate value theorem, which asserts that given the endpoints $R_0(\sigma_1, \sigma_2) < a$ and $R_1(\sigma_1, \sigma_2) = 1$, $\exists t' \in (0, 1)$, such that $R_{t'}(\sigma_1, \sigma_2) \in [a, b]$ (with high probability) because of the continuity of R_t .

6.2 Proof of Theorem 4.5

We now sketch a complete proof of [Theorem 4.5](#).

Proof. The input problem Ψ is a (k, d) -CSP(f) over n variables chose as in [Definition 1.4](#). This allows one to encode the problem instance Ψ as a random hypergraph $G \sim \mathcal{H}_{d,k,n}$ by choosing $h_j = f(x_{j_1}, \dots, x_{j_k})$ to be the local cost function for the j -th clause of Ψ . The collection of variables in each clause corresponds to the variables in a hyperedge in G_Ψ . Let the number of clauses be $|E|$. By definition, $|E| \sim \text{Poisson}(dn/k)$. Now, we set the hamiltonian as,

$$H_{d,k,n}^{G_\Psi} = \sum_{j=1}^{|E|} f(x_{j_1}, \dots, x_{j_k}).$$

By [Lemma 2.20](#), with the choice of p stated in the hypothesis, the $2p$ neighborhood will be no larger than n^A with probability $\geq 1 - e^{-n^a}$. The proof then concludes by invoking [Corollary 5.13](#) with $H_{d,k,n}^{G_\Psi}$, conditioned on the fact that the largest $2p$ neighborhood of $H_{d,k,n}^{G_\Psi}$ has size no more than n^A . $\square$

6.3 Proof of Theorem 4.3

We now sketch a complete proof of [Theorem 4.3](#).

Proof. Note that, once again, Ψ can be encoded into its representative hypergraph G_Ψ exactly as in the proof of [Theorem 4.5](#). Furthermore, by the hypothesis, the underlying problem (k, d) -CSP(f) satisfies a coupled OGP which can be interpreted as follows: For any $t \in [0, 1]$, choose $\text{Poisson}(tdn/k)$

clauses independently and uniformly at random and create an instance Ψ from them. Then, independently sample $\text{Poisson}((1-t)dn/k)$ clauses uniformly at random twice and create instances Ψ_1 and Ψ_2 from them. Let the final instances be $\Psi = \Psi_0 \cup \Psi_1$ and $\Psi' = \Psi_0 \cup \Psi_2$. This is exactly equivalent to the interpolation [Definition 2.18](#) over the representative hypergraphs G_Ψ and $G_{\Psi'}$. Note that the hypothesis implies we can assume these hypergraphs have $2p$ -neighborhoods of size no more than n^A with high probability ([Lemma 2.20](#)).

Concentration of objective value. Set the representative hamiltonians $H_{d,k,n}^{G_\Psi}$ and $H_{d,k,n}^{G_{\Psi'}}$ to be the sum of f acting on their respective clauses $\Psi_0 \cup \Psi_1$ and $\Psi_0 \cup \Psi_2$. Then, by [Theorem 4.5](#), both these functions will be concentrated with probability $\geq 1 - e^{-O(n^\gamma)}$.

Concentration of Hamming Weight. By invoking [Corollary 5.10](#) for each instance, we conclude that the hamming weight is concentrated with near certainty.

Concentration of Overlap of Output. This follows directly by applying [Theorem 5.1](#) to G_Ψ and $G_{\Psi'}$.

Given these properties, the argument in the proof for [Theorem 4.2](#) immediately implies the desired obstruction. $\square$

7 Strengthened McDiarmid's Inequality for biased distributions

We introduce the notions of a martingale and a Doob martingale, followed by a concentration result for bounded difference martingales in Fan et al. [[FGL12](#)], and conclude with a proof of a strengthened version of McDiarmid's inequality for highly-biased distributions.

7.1 Martingales & concentration

Definition 7.1 (Martingales).

A martingale with respect to random variables $X_1, \dots, X_n$ is given by a sequence of random variables $Z_1, \dots, Z_n$ such that the conditional expectation of each Z_i conditioned on all previous data points $X_{<i}$ in the sequence is,

$$\mathbb{E}[Z_i \mid X_1, \dots, X_{i-1}] = Z_{i-1}.$$

Furthermore, all expectations are bounded: $\mathbb{E}[|Z_i|] < \infty$. If $X_i = Z_i$, then $Z_1, \dots, Z_n$ is said to be a martingale with respect to itself.

A common martingale is the so-called *Doob martingale* where we set the random variables Z_i to be the averages of some bounded function f acting on $X_1, \dots, X_n$ conditioned on observations up to $X_{<i}$.

Definition 7.2 (Doob martingales).

The Doob martingale of a function f of random variables $X_1, \dots, X_n$ is the martingale sequence given by,

$$Z_i = \mathbb{E}[f(X_1, \dots, X_n) \mid X_1, \dots, X_i],$$

so long as $\mathbb{E}[|f(X_1, \dots, X_n)|] < \infty$.

We introduce some more technical definitions below.

Definition 7.3 (Martingale difference).

The martingale difference sequence of a martingale $Z_1, \dots, Z_n$ is given by the sequence $Y_i = Z_i - Z_{i-1}$, noting that $\mathbb{E}[Y_i] = 0$ by definition.

Definition 7.4 (Quadratic characteristic sequence).

The quadratic characteristic sequence of a martingale $Z_1, \dots, Z_n$ with martingale difference sequence $Y_2, \dots, Y_n$ consists of the values

$$\langle Z \rangle_i = \sum_{j \in [i]} \mathbb{E}[Y_j^2 \mid Z_1, \dots, Z_{j-1}].$$

To prove the version of McDiarmid's inequality with a highly-biased distribution, we will need a result that generalizes the concentration bounds of Azuma and Bernstein, using the quadratic characteristic sequence as the martingale analogue of the variance. The generalization is given by Fan et al. [FGL12] and was previously used to bound deviations of functions of the indicator variables of edges in sparse random graphs in Chen et al. [CST22].

Theorem 7.5 (Theorem 2.1 and Remark 2.1 combined with equation (11) of [FGL12]).

Let $Z_1, \dots, Z_n$ be a martingale with martingale differences Y_i satisfying $|Y_i| \leq 1$ for all $2 \leq i \leq n$. For every $0 \leq x \leq n$ and $v \geq 0$, we have

$$\Pr[|Z_n - Z_0| \geq x \text{ and } \langle Z \rangle_n \leq v^2] \leq 2 \exp\left(\frac{-x^2}{2(v^2 + x/3)}\right).$$

7.2 Strengthened McDiarmid's inequality

Lemma 7.6 (McDiarmid's inequality for biased distributions).

Suppose that $X_1, \dots, X_n$ are sampled i.i.d. from a distribution D over a finite set $\mathcal{X}$, such that D assigns probability $1 - p$ to a particular outcome $\chi_0 \in \mathcal{X}$. Let $f : \mathcal{X}^n \rightarrow \mathbb{R}$ satisfy a bounded-differences inequality, so that

$$|f(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n) - f(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)| \leq c$$

for all $x_1, \dots, x_n, x'_i \in \mathcal{X}$. Then

$$\Pr[|f(X_1, \dots, X_n) - \mathbb{E} f(X_1, \dots, X_n)| \geq \epsilon] \leq 2 \exp\left(\frac{-\epsilon^2}{2np(2-p)c^2 + 2c\epsilon/3}\right).$$

Proof. f is bounded because it has a finite domain. Therefore, let $Z_i = \mathbb{E}[f(X_1, \dots, X_n) \mid X_1, \dots, X_i]$ be the Doob martingale of $f(X_1, \dots, X_n)$ and let $Y_i = Z_i - Z_{i-1}$ be its martingale difference sequence.

We will show that $|Y_i| \leq c$. We take

$$g_j(x) := \mathbb{E}_{x_{j+1}, \dots, x_n \sim D} [f(X_1, \dots, X_{j-1}, x, x_{j+1}, \dots, x_n)],$$

so that by definition,

$$\begin{aligned} Y_i &= \mathbb{E}[f(X_1, \dots, X_n) \mid X_1, \dots, X_i] - \mathbb{E}[f(X_1, \dots, X_n) \mid X_1, \dots, X_{i-1}]. \\ &= \mathbb{E}_{x_{i+1}, \dots, x_n \sim D} f(X_1, \dots, X_{i-1}, X_i, x_{i+1}, \dots, x_n) - \mathbb{E}_{x_i, \dots, x_n \sim D} [f(X_1, \dots, X_{i-1}, x_i, x_{i+1}, \dots, x_n)] \\ &= g_i(X_i) - \mathbb{E}_{x \sim D} [g_i(x)] \\ &= \mathbb{E}_{x \sim D} [g_i(X_i) - g_i(x)]. \end{aligned}$$

Then it becomes a simple matter to compute

$$|Y_i| = \left| \mathbb{E}_{x \sim D} [g_i(X_i) - g_i(x)] \right| \leq \mathbb{E}_{x \sim D} [|g_i(X_i) - g_i(x)|] \leq c.$$

Next, we show that $\langle Z \rangle_i \leq (2-p)pc^2n$. By definition, $\langle Z \rangle_i = \sum_{j \in [i]} \mathbb{E}_{X_j \sim D} [Y_j^2 \mid X_1, \dots, X_{j-1}]$ and it will suffice to show that $\mathbb{E}_{X_j \sim D} [Y_j^2 \mid X_1, \dots, X_{j-1}] \geq (2-p)pc^2$. We start with

$$\mathbb{E}_{X_j \sim D} [Y_j^2] = \mathbb{E}_{X_j \sim D} \left[\left(\mathbb{E}_{x \sim D} g(X_j) - g(x) \right)^2 \right]$$

We split the expectation of Y_j^2 over $X_j \in \mathcal{X}$ into two cases, one where $X_j = \chi_0$ and one where $X_j \neq \chi_0$. Then

$$\mathbb{E}_{X_j \sim D} [Y_j^2] = \Pr_{X_j \sim D} [X_j = \chi_0] \mathbb{E}_{X_j \sim D} [Y_j^2 \mid X_j = \chi_0] + \Pr_{X_j \sim D} [X_j \neq \chi_0] \mathbb{E}_{X_j \sim D} [Y_j^2 \mid X_j \neq \chi_0]$$

Since $\Pr_{X_j \sim D} [X_j \neq \chi_0] = p$ and $|Y_i| \leq c$, we bound the second term in the above by pc^2 . Continuing with the first term,

$$\begin{aligned} \mathbb{E}_{X_j \sim D} [Y_j^2] &\leq pc^2 + \Pr_{X_j \sim D} [X_j = \chi_0] \left(g(\chi_0) - \mathbb{E}_{x \sim D} g(x) \right)^2 \\ &= pc^2 + \Pr_{X_j \sim D} [X_j = \chi_0] \left(g(\chi_0) - \Pr_{x \sim D} [x = \chi_0] g(\chi_0) - \Pr_{x \sim D} [x \neq \chi_0] \mathbb{E}_{x \sim D} [g(x) \mid x \neq \chi_0] \right)^2 \\ &= pc^2 + \Pr_{X_j \sim D} [X_j = \chi_0] \left(\Pr_{x \sim D} [x \neq \chi_0] g(\chi_0) - \Pr_{x \sim D} [x \neq \chi_0] \mathbb{E}_{x \sim D} [g(x) \mid x \neq \chi_0] \right)^2 \\ &= pc^2 + \Pr_{X_j \sim D} [X_j = \chi_0] \Pr_{x \sim D} [x \neq \chi_0] \left(g(\chi_0) - \mathbb{E}_{x \sim D} [g(x) \mid x \neq \chi_0] \right)^2 \\ &\leq pc^2 + (1-p)pc^2. \end{aligned}$$

If we scale down the martingale $Z_1, \dots, Z_n$ by a factor of c , we obtain a new martingale with martingale differences bounded by 1 in absolute value and with i th quadratic characteristic at most $p(2-p)i$. Therefore, by [Theorem 7.5](#), taking $v^2 = p(2-p)n$,

$$\Pr \left[\frac{1}{c} |Z_n - Z_0| \geq \epsilon \text{ and } \langle Z \rangle_n \leq p(2-p)n \right] \leq 2 \exp \left(\frac{-\epsilon^2}{2(2p(1-p)n + \epsilon/3)} \right).$$

Since $\langle Z \rangle_n \leq p(2-p)n$ is always true, and absorbing a factor of c into a rescaling of ϵ ,

$$\Pr [|Z_n - Z_0| \geq \epsilon] \leq 2 \exp \left(\frac{-\epsilon^2}{2np(2-p)c^2 + 2c\epsilon/3} \right).$$

□

8 Overlap-Gap Property for general case random Max- k -XOR

In this section we prove that the coupled OGP exists even for signed versions of the random Max- k -XOR problem by extending the proof from [\[CGP⁺19, Section 4\]](#).

8.1 The signed Max- k -XOR hamiltonian

We extend [\[CGP⁺19, Lemma 4.1\]](#) to handle an interpolation with *random* signs for the variables. Explicitly, we will consider the following hamiltonian,

$$H_{\text{signed}}^G = - \sum_{i=1}^m \prod_{j=1}^k p_{ij} \sigma_{v_{ij}}, \quad (8.1)$$

where $p_{ij} \sim \{\pm 1\}$ are i.i.d. Rademacher random variables ([Definition 2.4](#)). Notice that when we deterministically fix $p_{ij} = 1$ for every j -th variable in the i -th clause, this recovers the diluted k -spin glass hamiltonian H^G . Maximizing H_{signed}^G corresponds to solving a random instance of a *signed* Max- k -XOR problem.

On expectation under the Rademacher distribution, the random signs will be balanced. This property allows us to amend the Poisson integration techniques in [\[CGP⁺19, Lemma 4.1\]](#) to preserve the asymptotic behavior of the hamiltonian to be equivalent, up to a constant shift, to the unsigned setting under the coupled Guerra-Toninelli interpolation [\[GT04\]](#).

8.2 The Guerra-Toninelli interpolations

We work with two families of interpolations: The first is a gaussian interpolation between independent copies of k -mean field hamiltonians and the second is an interpolation between a diluted spin-glass hamiltonian and a dense one. For the gaussian interpolation, choose an instance of a k -mean field hamiltonian H_k with independent copies H'_k and H''_k and interpolate smoothly as

$$H_k^1 = \sqrt{t}H_k + \sqrt{1-t}H'_k, \quad (8.2)$$

$$H_k^2 = \sqrt{t}H_k + \sqrt{1-t}H''_k, \quad (8.3)$$

where $t \in [0, 1]$. The diluted-to-dense interpolation is known as the Guerra-Toninelli interpolation [GT04] and is given as

$$H(s, \sigma^1, \sigma^2) = \sum_{l=1}^2 \left(\delta H_{\text{signed}, \frac{d}{k}(1-s), t}^l(\sigma^l) + \sqrt{s} \beta H_k^l(\sigma^l) \right), \quad (8.4)$$

where $H_{\text{signed}, \frac{d}{k}(1-s), t}^1$ and $H_{\text{signed}, \frac{d}{k}(1-s), t}^2$ are drawn from the distribution of the coupled interpolation defined in Definition 2.18 with $\text{Poisson}(\frac{d}{k}(1-s))$ edges and the additional requirement that the Rademacher variables also be re-sampled for every $t \in [0, 1]$.

The last notion we need is that of an average with respect to the so-called *Gibbs Measure*, which is a normalized probability given to every pair of configurations weighted by their coupled energy at time s . The Gibbs measure over $A \subseteq \{\pm 1\}^n$ and overlap set $S \subseteq [0, 1]$ is defined as

$$G_s(\sigma^1, \sigma^2) = \frac{\exp(H(s, \sigma^1, \sigma^2))}{\sum_{\sigma^1, \sigma^2 \in A, |R_{1,2}| \in S} \exp(H(s, \sigma^1, \sigma^2))}. \quad (8.5)$$

An average of a quantity $\mathcal{B}$ with respect to the Gibbs measure is denoted as $\langle \mathcal{B} \rangle_s$. The denominator of Equation 8.5, denoted as Z , is a normalization term called the *partition function*.

8.3 Coupled Overlap-Gap Property for general case Max- k -XOR

Lemma 8.6 (Scaling of random Max- k -XOR under signed Guerra-Toninelli interpolation).

For any $A \subseteq \{\pm 1\}^n$, $S \subseteq [0, 1]$, $t \in [0, 1]$ and $d, k > 0$, the following holds,

$$\begin{aligned} & \frac{1}{n} \mathbb{E} \left[\max_{\sigma^1, \sigma^2 \in A, |R_{1,2}| \in S} H_{\text{signed}}^1(\sigma^1) + H_{\text{signed}}^2(\sigma) \right] \\ &= \frac{1}{n} \sqrt{\frac{d}{k}} \mathbb{E} \left[\max_{\sigma^1, \sigma^2 \in A, |R_{1,2}| \in S} H_k^1(\sigma^1) + H_k^2(\sigma) \right] + O \left(\left(\frac{d}{k} \right)^{1/3} \right). \end{aligned} \quad (8.7)$$

Proof. We split the H_{signed}^G terms in the Guerra-Toninelli interpolated hamiltonian into two terms. In the first term, we will collect the positive signed hyperedges (Π^+) and in the second term we will collect the negative signed hyperedges (Π^-).

$$\begin{aligned} H_{\text{signed}}^G &= - \sum_{i=1}^m \prod_{j=1}^k p_{ij} \sigma_{v_{ij}} \\ &= - \sum_{i=1}^m \prod_{j=1}^k (p_{ij}) \prod_{l=1}^k (\sigma_{v_{ij}}) \\ &:= - \sum_{i=1}^m p_i \prod_{j=1}^k (\sigma_{v_{ij}}), \end{aligned}$$

where $p_i = \prod_{j=1}^k (p_{ij})$ are i.i.d. Rademacher random variables for $i \in [m]$. This follows because each p_{ij} is an independent Rademacher variable, and therefore, $\prod_{j=1}^k p_{ij} \sim \{\pm 1\}$ is also a Rademacher variable. The sum is now split into positive terms ($p_{ij} = 1$) and negative terms ($p_{ij} = -1$).

In other words, we may write

$$H_{\text{signed}, \frac{d}{k}(1-s), t}^\ell = H_{\frac{d}{k}(1-s)/2, t}^\ell - H_{\frac{d}{k}(1-s)/2, t}^{\ell'},$$

where $H_{\frac{d}{k}(1-s)/2, t}^\ell$ and $H_{\frac{d}{k}(1-s)/2, t}^{\ell'}$ are two independent copies of the coupled distribution defined in [Definition 2.18](#) with $\text{Poisson}(\frac{d}{k}(1-s)/2)$ hyperedges. So in the interpolation, we write

$$\phi(s) = \frac{1}{n} \mathbb{E} \left[\log \sum \exp \left[\sum_{\ell \in \{1,2\}} \sqrt{s} \beta H_k^\ell(\sigma^\ell) + \delta H_{\frac{d}{k}(1-s)/2, t}^\ell(\sigma^\ell) - \delta H_{\frac{d}{k}(1-s)/2, t}^{\ell'}(\sigma^\ell) \right] \right]$$

We break the interpolation into a sum of a Gaussian interpolation and two Poisson interpolations,

$$\phi'(s) = \text{I} + \text{II}^+ + \text{II}^-,$$

so that

$$\text{I} = \mathbb{E} \left[\frac{1}{Z} \left(\sum \exp(H(s, \sigma^1, \sigma^2)) \sum_{\ell \in \{1,2\}} \frac{1}{2\sqrt{s}} \beta H_k^\ell(\sigma^\ell) \right) \right] = \sum_{\ell \in \{1,2\}} \left\langle \mathbb{E} \left[\frac{1}{2\sqrt{s}} \beta H_k^\ell(\sigma^\ell) \right] \right\rangle_s,$$

for which we may then apply Stein's Lemma. A similar calculation applies for II^+ and II^- , for which we use the fact about Poisson random variables that

$$\frac{d}{ds} \mathbb{E} f(\text{Poisson}(s)) = \mathbb{E} f(\text{Poisson}(s) + 1) - \mathbb{E} f(\text{Poisson}(s)),$$

as in [\[CGP⁺19, Proof of Lemma 4.1, Page 13\]](#).

Case I: Positive p_{ij} . This is equivalent to the terms obtained by [\[CGP⁺19, Proof of Lemma 4.1, Page 15\]](#). Namely, the following term is obtained (after Taylor Expansion),

$$\text{II}^+ = -\frac{d}{k} \log \cosh(\delta) + \frac{d}{2k} \sum_{r \geq 1} \frac{\tanh(\delta)^r}{r} \left(t \mathbb{E}[\langle \Delta(\sigma^1, \sigma^2) \rangle_s^r] + (1-t) \sum_{l=1}^2 \mathbb{E}[\langle \sigma_{i_1}^l \cdots \sigma_{i_k}^l \rangle_s^r] \right)$$

Using replicas to represent the difference term $\Delta(\sigma^1, \sigma^2)$ and then evaluating it under a random choice of indices, followed by a second-order Taylor expansion of the preceding term, [\[CGP⁺19\]](#) obtain

$$\begin{aligned} \text{II}^+ &= -\frac{d}{k} \log \cosh(\delta) + \frac{d}{2k} \tanh(\delta) \mathbb{E}[\langle m(\sigma^{1,1})^k + m(\sigma^{1,2})^k \rangle_s] - t \frac{d}{2k} \tanh(\delta)^2 \mathbb{E}[\langle (R_{1,2}^{1,1})^k \rangle_s] \\ &\quad + \frac{d}{2k} \frac{\tanh(\delta)^2}{2} \mathbb{E}[\langle (R_{1,2}^{1,1})^k \rangle_s + \langle (R_{2,2}^{1,2})^k \rangle_s + t \left(\langle (R_{1,2}^{1,2})^k \rangle_s + \langle (R_{2,1}^{1,2})^k \rangle_s \right)] + O\left(\frac{d}{k} \delta^3\right). \end{aligned}$$

We divide the term in [\[CGP⁺19, Lemma 4.1\]](#) by $\frac{1}{2}$ since we are working with half the edges (in expectation) in the modified interpolation.

Case II: Negative p_{ij} . This case is equivalent to that of positive signs up to a change in the sign of a field term that depends on the overlap between two replicas. This sign flip eliminates the

magnetization term that appears in the statement of [CGP⁺19, Lemma 4.1]. To this extent, we define a modified Δ function called Δ^- , similar to [CGP⁺19].

$$\begin{aligned}\Pi^- &= -\frac{dt}{2k} \left(\mathbb{E} \log \sum \exp H^-(s, \sigma^1, \sigma^2) - \mathbb{E} \log \sum \exp H(s, \sigma^1, \sigma^2) \right) \\ &\quad - \frac{(1-t)d}{2k} \left(\mathbb{E} \log \sum \exp H_1^-(s, \sigma^1, \sigma^2) - \mathbb{E} \log \sum \exp H(s, \sigma^1, \sigma^2) \right) \\ &\quad - \frac{(1-t)d}{2k} \left(\mathbb{E} \log \sum \exp H_2^-(s, \sigma^1, \sigma^2) - \mathbb{E} \log \sum \exp H(s, \sigma^1, \sigma^2) \right) \\ \Delta^-(\sigma^1, \sigma^2) &= \sigma_{j_1}^1 \cdots \sigma_{j_K}^1 + \sigma_{j_1}^2 \cdots \sigma_{j_K}^2 + \tanh(\delta) \sigma_{j_1}^1 \cdots \sigma_{j_K}^1 \sigma_{j_1}^2 \cdots \sigma_{j_K}^2\end{aligned}$$

Repeating the same perturbation based calculation on the interpolated hamiltonian with $p_e = -1$ in front of every hyperedge e as is done in the positive case, it is not hard to obtain

$$\begin{aligned}\Pi^- &= -\frac{d}{k} \log \cosh \delta \\ &\quad - \frac{d}{2k} \sum_{r=1}^{\infty} (-1)^{r-1} \frac{\tanh(\delta)^r}{r} \left(t \mathbb{E} \left\langle \Delta^-(\sigma^1, \sigma^2) \right\rangle_s^r + (1-t) \left[\mathbb{E} \left\langle \sigma_{j_1}^1 \cdots \sigma_{j_K}^1 \right\rangle_s^r + \mathbb{E} \left\langle \sigma_{j_1}^2 \cdots \sigma_{j_K}^2 \right\rangle_s^r \right] \right)\end{aligned}$$

By the arguments in the proof of [CGP⁺19, Lemma 4.1], introducing replicas $\sigma^{\ell,1}$ and $\sigma^{\ell,2}$ for σ^1 and σ^2 , the first term of this Taylor series is

$$\begin{aligned}&\tanh(\delta) \left(t \mathbb{E} \left\langle \Delta^-(\sigma^1, \sigma^2) \right\rangle_s + (1-t) \mathbb{E} \left\langle \sigma_{j_1}^1 \cdots \sigma_{j_K}^1 \right\rangle_s + (1-t) \mathbb{E} \left\langle \sigma_{j_1}^2 \cdots \sigma_{j_K}^2 \right\rangle_s \right) \\ &= \tanh(\delta) \left(t \mathbb{E} \left\langle m(\sigma^{1,1})^k + m(\sigma^{1,2})^k + \tanh(\delta) (R_{1,2}^{1,1})^k \right\rangle_s + (1-t) \mathbb{E} \left\langle m(\sigma^{1,1})^k + m(\sigma^{1,2})^k \right\rangle_s \right) \\ &= \tanh(\delta) \left(\mathbb{E} \left\langle m(\sigma^{1,1})^k + m(\sigma^{1,2})^k \right\rangle_s + t \tanh(\delta) \mathbb{E} \left\langle (R_{1,2}^{1,1})^k \right\rangle_s \right)\end{aligned}$$

We now compute the second-order term of the Taylor expansion above and evaluate the terms $\mathbb{E}[\langle \Delta^-(\sigma^1, \sigma^2) \rangle_s^2]$ and $\sum_{l=1}^2 \mathbb{E}[\langle \sigma_{j_1}^l \cdots \sigma_{j_K}^l \rangle_s^2]$.

Using the introduced replicas and averaging over the indices of the hyperedges as in [CGP⁺19, Proof of Lemma 4.1] yields

$$\begin{aligned}\mathbb{E}[\langle \Delta^-(\sigma^1, \sigma^2) \rangle_s^2] &= \mathbb{E} \left[\left\langle \mathbb{E}' \left[\left(\sum_{l=1}^2 \prod_{r=1}^k \sigma_{j_r}^{1,l} + \tanh(\delta) (R_{1,2}^{1,1})^k \right) \left(\sum_{l=1}^2 \prod_{r=1}^k \sigma_{j_r}^{2,l} + \tanh(\delta) (R_{1,2}^{2,2})^k \right) \right] \right\rangle_s \right] \\ &= \mathbb{E} \left[\left\langle \sum_{r_1, r_2=1}^2 (R_{r_1, r_2}^{1,2})^k \right\rangle_s + \tanh(\delta) \left((R_{1,2}^{2,2})^k \sum_{l=1}^2 \prod_{r=1}^k \sigma_{j_r}^{1,l} + (R_{1,2}^{1,1})^k \sum_{l=1}^2 \prod_{r=1}^k \sigma_{j_r}^{2,l} \right) + \mathcal{O}(\tanh(\delta)^2) \right].\end{aligned}\tag{8.8}$$

As in [CGP⁺19, Proof of Lemma 4.1, Pg 15], for the second term we have that

$$\sum_{l=1}^2 \mathbb{E}[\langle \sigma_{i_1}^l \cdots \sigma_{i_K}^l \rangle_s^2] = \mathbb{E} \left[\left\langle \sum_{r=1}^2 (R_{r,r}^{1,2})^k \right\rangle_s \right].$$

Using the facts that the $\Delta^-(\sigma^1, \sigma^2) \leq 3$ and $|\sigma_{i_1}^1 \cdots \sigma_{i_k}^1| \leq 1$, all terms of order ≥ 3 in the Taylor expansion are no more than $L \frac{d}{k} \delta^3$ for an appropriate $L > 0$. Putting together the expansions yields the following bound,

$$\begin{aligned} \Pi^- = & -\frac{d}{k} \log \cosh(\delta) - \frac{d}{2k} \tanh(\delta) \left(\mathbb{E} \left\langle m(\sigma^{1,1})^k + m(\sigma^{1,2})^k \right\rangle_s + t \tanh(\delta) \mathbb{E} \left\langle (R_{1,2}^{1,1})^k \right\rangle_s \right) \\ & + \frac{td}{2k} \frac{\tanh(\delta)^2}{2} \left(\mathbb{E} \left[\left\langle \sum_{r_1, r_2=1}^2 (R_{r_1, r_2}^{1,2})^k \right\rangle_s + \tanh(\delta) \left((R_{1,2}^{2,2})^k \sum_{l=1}^2 \prod_{r=1}^k \sigma_{jr}^{1,l} + (R_{1,2}^{1,1})^k \sum_{l=1}^2 \prod_{r=1}^k \sigma_{jr}^{2,l} \right) + \mathcal{O}(\tanh(\delta)^2) \right] \right) \\ & + \frac{d(1-t)}{2k} \frac{\tanh(\delta)^2}{2} \left(\mathbb{E} \left[\left\langle \sum_{r=1}^2 (R_{r,r}^{1,2})^k \right\rangle_s \right] \right). \end{aligned}$$

Notice that all terms of $\mathcal{O}(\tanh(\delta)^3)$ and higher powers therein can be absorbed into the term $L \frac{d}{k} \delta^3$. This finally yields

$$\begin{aligned} \Pi^- = & -\frac{d}{k} \log \cosh(\delta) - \frac{d}{2k} \tanh(\delta) \left(\mathbb{E} \left\langle m(\sigma^{1,1})^k + m(\sigma^{1,2})^k \right\rangle_s \right) - \frac{td}{2k} \tanh(\delta)^2 \left(\mathbb{E} \left\langle (R_{1,2}^{1,1})^k \right\rangle_s \right) \\ & + \frac{d}{2k} \frac{\tanh(\delta)^2}{2} \left(\mathbb{E} \left[\left\langle t (R_{1,2}^{1,2})^k + t (R_{2,1}^{1,2})^k + (R_{1,1}^{1,2})^k + (R_{2,2}^{1,2})^k \right\rangle_s \right] \right) + \mathcal{O} \left(\frac{d}{2k} \delta^3 \right). \quad (8.9) \end{aligned}$$

Combining Π^+ and Π^- . We now add the positive and negative terms together to obtain the equivalent of Π for H_{signed}^G . As a result of the sign flip, the magnetization dependencies cancel out. This finally yields

$$\begin{aligned} \Pi = \Pi^+ + \Pi^- = & -\frac{2d}{k} \log \cosh(\delta) - \frac{td}{k} \tanh(\delta)^2 \left(\mathbb{E} \left\langle (R_{1,2}^{1,1})^k \right\rangle_s \right) \\ & + \frac{d}{k} \frac{\tanh(\delta)^2}{2} \left(\mathbb{E} \left[\left\langle t (R_{1,2}^{1,2})^k + t (R_{2,1}^{1,2})^k + (R_{1,1}^{1,2})^k + (R_{2,2}^{1,2})^k \right\rangle_s \right] \right) + \mathcal{O} \left(\frac{d}{k} \delta^3 \right). \quad (8.10) \end{aligned}$$

For fixed d, k and δ , we define β as

$$\beta = \sqrt{\frac{d}{k}} \tanh(\delta).$$

This choice causes the overlap terms to cancel in $\phi(s)$, yielding the following rate of change of the free energy

$$\phi'(s) = \text{I} + \text{II} = -\frac{d}{k} \log \cosh(\delta) + \frac{d}{k} \tanh(\delta)^2 + \mathcal{O} \left(\frac{d}{k} \delta^3 \right) \xrightarrow{\delta \rightarrow 0} \mathcal{O} \left(\frac{d}{k} \delta^3 \right) + \mathcal{O} \left(\frac{d}{k} \delta^4 \right), \quad (8.11)$$

where I is defined equivalently as in the proof of [CGP⁺19, Lemma 4.1, Pg 13]. This immediately yields that

$$\frac{1}{\delta} (\phi(1) - \phi(0)) = \frac{1}{\delta} \int_0^1 \phi'(s) ds = \mathcal{O} \left(\frac{d}{k} \delta^2 \right).$$

The rest of the argument follows exactly as in [CGP⁺19, Proof of Lemma 4.1, Pg 16] with the final substitution $\delta = \left(\frac{d}{k} \right)^{-1/3}$. $\square$

Having proved the key interpolation lemma about the coupled free energies in the diluted and dense models, we now state the coupled OGP for the random *signed* Max- k -XOR problem.

Theorem 8.12 (Coupled OGP for random Max- k -XOR, k even).

For every even $k \geq 4$, there exists an interval $0 < a < b < 1$ and parameters $d_0 > 0$, $0 < \eta_0 < P(k)$ and $n_0 > 1$, such that, for any $t \in [0, 1]$, $d \geq d_0$, $n \geq n_0$ and constant $L = L(\eta_0, d)$, with probability at least $1 - Le^{-n/L}$ over the t -coupled Max- k -XOR instance pair $(\Psi_1, \Psi_2) \sim \mathcal{H}_{n,d,k,t}^\pm$, whenever two spins σ_1, σ_2 satisfy

$$\frac{H^{\Psi_i}(\sigma_i)}{n} \geq M(k, d) \left(1 - \frac{\eta_0}{P(k)}\right),$$

then their overlap satisfies $|R(\sigma_1, \sigma_2)| \notin [a, b]$.

Proof. This can be proven by an exact copy of [CGP⁺19, Theorem 5], but using Lemma 8.6 instead of [CGP⁺19, Lemma 4.1]. $\square$

To obstruct local algorithms using the same framework mentioned in Section 1.4, it is critical that the overlap between nearly optimal solutions of independent random Max- k -XOR instances (with signs) be small. To show this, we first prove that if the overlap between pairs of solutions of two independent instances of the k -mean field model is bounded away from 0, then they are suboptimal.

Lemma 8.13.

Consider the parameters $[a, b]$ and η_0 from Theorem 8.12. For large enough n , there is $\hat{\eta} > 0$ satisfying $\hat{\eta} > \eta_0$ such that

$$\frac{1}{n} \mathbb{E} \left[\max_{|R_{1,2}| \in [a, 1]} (H_k^1(\sigma^1) + H_k^2(\sigma^2)) \right] < 2(P(k) - \hat{\eta})$$

where H_k^1 and H_k^2 are random 0-coupled instances of a k -spin glass.

Proof. In the proof of Theorem 8.12 as in that of [CGP⁺19, Theorem 5], the parameter a is chosen from [CGP⁺19, Theorem 3], and the parameter η_0 is chosen so that

$$\limsup_{n \rightarrow \infty} \mathbb{E} \frac{1}{n} \max_{|R_{1,2}| \in [a, b]} (H_k^1(\sigma^1) + H_k^2(\sigma^2)) < 2P(k) - 6\eta_0,$$

for all $t \in [0, 1]$, where $R_{1,2}$ is the overlap between σ^1 and σ^2 , and H_k^1 and H_k^2 are t -coupled instances of a k -spin glass.

However, in fact for the $t = 0$ case (as noted in the remark in the proof of [CGP⁺19, Theorem 5]), the proof cites [CP18, Theorem 2], which provides a bound for $|R_{1,2}| \in [a, 1]$ with $|R_{1,2}| \in [a, b]$ being a subcase of that. This implies

$$\limsup_{n \rightarrow \infty} \mathbb{E} \frac{1}{n} \max_{|R_{1,2}| \in [a, 1]} (H_k^1(\sigma^1) + H_k^2(\sigma^2)) < 2P(k) - 6\eta_0,$$

at $t = 0$.

Continuing in the special case where $t = 0$, then, let η be so that

$$\limsup_{n \rightarrow \infty} \mathbb{E} \frac{1}{n} \max_{|R_{1,2}| \in [a,1]} (H_k^1(\sigma^1) + H_k^2(\sigma^2)) = 2(P(k) - \eta),$$

Since $\eta \geq 3\eta_0$, we choose $\hat{\eta} = \eta + \epsilon$, for some ϵ which is allowed to be arbitrarily small as $n \rightarrow \infty$. □

We now extend [Lemma 8.13](#) to setting of the diluted model with signs, and show that pairs of nearly optimal solutions of independent random Max- k -XOR instances (signed) have low overlap.

Lemma 8.14.

Consider the parameters $[a, b]$ and η_0 from [Theorem 8.12](#). For large enough n , there are $\eta', L > 0$ with $\eta' > \eta_0$ such that

$$\begin{aligned} \frac{1}{n} \max_{\substack{\sigma^1, \sigma^2 \in \{\pm 1\}^n, \\ |R_{1,2}| \geq a}} (H_{\text{signed}}^1(\sigma^1) + H_{\text{signed}}^2(\sigma^2)) &\leq 2(P(k) - \eta') \sqrt{\frac{d}{k}} + O(\sqrt[3]{d/k}) \\ &\leq 2M(k, d) \left(1 - \frac{\eta'}{P(k)} + O\left(\left(\frac{d}{k}\right)^{-1/6}\right) \right). \end{aligned}$$

with probability at least $1 - 2Le^{n/L}$ over the random choice of 0-coupled Max- k -XOR instances H_{signed}^1 and H_{signed}^2 .

Proof. By [Lemma 8.13](#), there is a choices of $\hat{\eta}$ satisfying $\hat{\eta} > \eta_0$ such that

$$\mathbb{E} \frac{1}{n} \max_{\substack{\sigma^1, \sigma^2 \in \{\pm 1\}^n \\ |R_{1,2}| \in [a,1]}} (H_k^1(\sigma^1) + H_k^2(\sigma^2)) < 2(P(k) - \hat{\eta}),$$

where H_k^1 and H_k^2 are still 0-coupled instances of the mean-field k -spin glass and $R_{1,2}$ is the overlap between σ^1 and σ^2 .

By instantiating [Lemma 8.6](#) with $S = [a, 1]$ and $A = \{\pm 1\}^n$ and $t = 0$, we see that

$$\frac{1}{n} \mathbb{E} \max_{\substack{\sigma^1, \sigma^2 \in \{\pm 1\}^n \\ |R_{1,2}| \in [a,1]}} (H_{\text{signed}}^1(\sigma^1) + H_{\text{signed}}^2(\sigma^2)) = \frac{1}{n} \sqrt{\frac{d}{k}} \mathbb{E} \max_{\substack{\sigma^1, \sigma^2 \in \{\pm 1\}^n \\ |R_{1,2}| \in [a,1]}} (H_k^1(\sigma^1) + H_k^2(\sigma^2)) \pm O(\sqrt[3]{d/k}),$$

Combining the two above inequalities,

$$\frac{1}{n} \mathbb{E} \max_{\substack{\sigma^1, \sigma^2 \in \{\pm 1\}^n \\ |R_{1,2}| \in [a,1]}} (H_{\text{signed}}^1(\sigma^1) + H_{\text{signed}}^2(\sigma^2)) \leq 2(P(k) - \hat{\eta}) \sqrt{\frac{d}{k}} + O(\sqrt[3]{d/k})$$

By an application of Azuma's inequality and concentration of Poisson random variables, there is some L as a function of $\hat{\eta}$ and η_0 such that

$$\frac{1}{n} \max_{\substack{\sigma^1, \sigma^2 \in \{\pm 1\}^n, \\ |R_{1,2}| \geq a}} (H_{\text{signed}}^1(\sigma^1) + H_{\text{signed}}^2(\sigma^2)) \leq 2(P(k) - \eta') \sqrt{\frac{d}{k}} + O(\sqrt[3]{d/k})$$

for some η' satisfying $\eta_0 < \eta' < \hat{\eta}$, with probability at least $1 - 2Le^{-n/L}$. $\square$

9 Discussion

Our work conclusively establishes the coupled OGP as an obstruction to all local quantum algorithms on *any* (k, d) -CSP(f). In doing this, the work hints at and leaves open many interesting questions for future work in areas that are at the intersection of Quantum inapproximability, Statistical Physics, Random Graph Theory, Combinatorial Optimization and Average-Case Complexity.

9.1 Which CSPs have an OGP?

While various sparse CSPs such as k -SAT, unsigned max- k -XOR and k -NAE-SAT have been shown to exhibit clustering in their solution spaces at different clause-to-variable ratios [ART06, DSS16, CGP⁺19], it is not known whether this property is pervasive to most CSPs or something that happens to a select few. Therefore, in order to understand the complexity landscape of CSPs on typical instances better, the following open question is interesting to investigate:

Conjecture 9.1 (Random Predicate CSPs and coupled OGP).

Given a function f chosen uniformly at random from the set of functions $\mathcal{B}_k = \{g \mid g : \{\pm 1\}^k \rightarrow \{0, 1\}\}$, (k, d) -CSP(f) has a coupled-OGP for sufficiently large k and d with high probability (over the choice of f and instance $\Psi \sim (k, d)$ -CSP(f)).

Notice that the conjecture above is specifically interested in the solution geometry of a CSP in the *unsatisfiable* regime (large d). A positive resolution to the above conjecture will make the obstructions stated in Theorem 1.6 hold for *almost all* CSPs for the family of generic local algorithms. Another question of interest is which properties (if any) about a predicate f can be identified which would conclusively imply that a random instance Ψ of a (k, d) -CSP(f) will have an OGP.

Problem 9.2 (Properties of coupled-OGP predicates).

Can we enumerate a set of necessary and sufficient conditions on f to be such that (k, d) -CSP(f) satisfies a coupled-OGP for sufficiently large k and d ?

9.2 Beyond log-depth obstructions for QAOA_p?

Work on obstructing QAOA_p using an OGP heavily relies on the locality of the algorithm at shallow depths. It is interesting to investigate whether this obstruction can be extended beyond the $\epsilon \log(n)$ -depth regime to make this a non-local obstruction. Recent work [GJW20, Wei22] suggests that the

OGP may actually result in stronger obstructions than just local ones, and it would be interesting to see if these techniques can be generalized to the setting of $QAOA_p$ to yield obstructions that are non-local.

Problem 9.3 (Polylogarithmic obstructions to $QAOA_p$ in the OGP regime).

Given a $QAOA_p$ circuit with depth $p \leq \epsilon (\log(n))^c$ for some $c > 1$, does there exist $\epsilon_0 > 0$, such that $QAOA_p$ is obstructed on a (k, d) -CSP(f) with a coupled OGP from outputting solutions that are better than $(1 - \epsilon_0)$ approximations to the optimal?

9.3 A Quantum OGP and lifting “classical” obstructions

The idea of the OGP obstructing families of algorithms that are *stable* under small perturbations [GJW20] motivates the idea of a quantized version of the OGP, to apply to *quantum* CSPs. To define such a property over quantum states, however, there would need to be a metric that is very similar to the classical hamming distance over $\mathbb{F}_2$ and has the property that it is invariant over permutations of the canonical basis, while still quantifying entanglement in a desired way. One such possible metric is a *quantum* version of the Wasserstein distance of first order that was proposed by De-Palma et al. [DPMTL21]. In particular, given a natural generalization of Definition 1.5 to a quantized setting using a quantum version of the Wasserstein distance of first order, it is interesting to investigate if a larger family of quantum circuits up to some depth $p(n)$ can be obstructed by a family of d -local hamiltonians $\{H_n\}_{n \geq n_0}$ that possess a qOGP (quantized Overlap-Gap Property). A result of this type could imply a way to generically “lift” classical obstructions for *stable* classical algorithms to a corresponding family of quantum algorithms.

9.4 Message-Passing algorithm for MAX-CUT of *all* d -regular graphs?

Finding an efficient classical algorithm that can output cuts that are arbitrary approximations of the optimal ones for d -regular graphs is a long-standing open problem in Random Graph Theory and Theoretical Computer Science. Recently, this problem was nearly completely solved by Alaoui et al. [AMS21] as they constructed a Message-Passing algorithm for random regular graphs of very large degree under the widely believed no-OGP assumption about the SK model. However, the problem does not provide a complete solution as it needs the degree d to be larger than $O(\frac{1}{\epsilon})$ in order to output a $(1 - \epsilon)$ -optimal cut. A natural question is whether, under a no-OGP assumption, the result can be extended to output $(1 - \epsilon)$ -optimal cuts for d -regular graphs for *any* $d \geq 3$.

Conjecture 9.4 (AMP algorithm for Random d -Regular Graphs).

There exists a $\text{poly}(n, \frac{1}{\epsilon})$ time algorithm A that outputs a $(1 - \epsilon)$ -approximate cut of a random d -regular graph G with high probability under a “no-OGP” assumption for any $d \geq 3$.

Note that the approach of Alaoui et al. [AMS21] critically relies on the Guerra-Tonnineli interpolation between the $\mathcal{G}_{n,d}$ model and the SK-model which will *only* work for $d \geq O(\frac{1}{\epsilon})$. Consequently, a solution that works for *all* $d \geq 3$ will require a fundamentally different approach. A natural question that is motivated by the above conjecture is to then investigate if there is any range of

degree for which the Max-Cut problem over d -regular graphs possesses an OGP. Given the belief that the SK model does not exhibit an OGP, this would only be an interesting question in the relatively low-degree regime.

Problem 9.5 (Random d -Regular Graphs don't have an OGP).

Does the MAX-CUT problem on random d -regular graphs have an OGP for some $d \geq 3$? If so, for what $\{d_0, d_1\} \subset \mathbb{N}$ does the problem exhibit an OGP?

The $QAOA_p$ algorithm was initiated and analyzed on the MAX-CUT of d -regular graphs and positive answers to the conjectures above will close the scope for any quantum advantage on the problem.

Acknowledgements

We thank Jonathan Wurtz for many insightful discussions about QAOA. We are grateful to Amartya Shankha Biswas for patiently explaining the *factors of i.i.d.* framework to us. We would also like to thank Antares Chen for many invigorating and profound discussions which culminated as the open problem proposed in [Problem 9.2](#). Lastly, we would like to thank Boaz Barak for providing detailed and helpful feedback on a prior version of this manuscript, and David Gamarnik for his explanations on the state of the art results in the research area.

References

- [AAB⁺20] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Sergio Boixo, Michael Broughton, Bob B. Buckley, David A. Buell, et al. Hartree-Fock on a superconducting qubit quantum computer. *Science*, 369(6507):1084–1089, 2020.
- [ACZ20] Antonio Auffinger, Wei-Kuo Chen, and Qiang Zeng. The SK model is infinite step replica symmetry breaking at zero temperature. *Communications on Pure and Applied Mathematics*, 73(5), 2020.
- [AMS21] Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Local algorithms for Maximum Cut and Minimum Bisection on locally treelike regular graphs of large degree. *arXiv preprint arXiv:2111.06813*, 2021.
- [ART06] Dimitris Achlioptas and Federico Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 130–139, 2006.
- [BBF⁺18] Fernando GSL Brandao, Michael Broughton, Edward Farhi, Sam Gutmann, and Hartmut Neven. For fixed control parameters the quantum approximate optimization algorithm’s objective function value concentrates for typical instances. *arXiv preprint arXiv:1812.04170*, 2018.
- [Bel64] John S Bell. On the Einstein-Podolsky-Rosen paradox. *Physics Physique Fizika*, 1(3):195, 1964.
- [BKKT20] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang. Obstacles to variational quantum optimization from symmetry protection. *Physical review letters*, 125(26):260505, 2020.
- [BM21] Boaz Barak and Kunal Marwaha. Classical algorithms and quantum limitations for maximum cut on high-girth graphs. *arXiv preprint arXiv:2106.05900*, 2021.
- [BMZ05] Alfredo Braunstein, Marc Mézard, and Riccardo Zecchina. Survey propagation: An algorithm for satisfiability. *Random Structures & Algorithms*, 27(2):201–226, 2005.

- [BPS99] Sally C Brailsford, Chris N Potts, and Barbara M Smith. Constraint satisfaction problems: Algorithms and applications. *European journal of operational research*, 119(3):557–581, 1999.
- [Can17] Clément Canonne. A short note on Poisson tail bounds. Retrieved from the website: <http://www.cs.columbia.edu/ccanonne>, 2017.
- [CGP⁺19] Wei-Kuo Chen, David Gamarnik, Dmitry Panchenko, Mustazee Rahman, et al. Sub-optimality of local algorithms for a class of max-cut problems. *Annals of Probability*, 47(3):1587–1618, 2019.
- [Cho94] K. P. Choi. On the Medians of Gamma Distributions and an Equation of Ramanujan. *Proceedings of the American Mathematical Society*, 121(1):245–251, 1994.
- [Con18] Francesco Concetti. The full replica symmetry breaking in the Ising spin glass on random regular graph. *Journal of Statistical Physics*, 173(5):1459–1483, 2018.
- [CP18] Wei-Kuo Chen and Dmitry Panchenko. Disorder chaos in some diluted spin glass models. *The Annals of Applied Probability*, 28(3):1356–1378, 2018.
- [CST22] Antares Chen, Jonathan Shi, and Luca Trevisan. Cut Sparsification of the Clique Beyond the Ramanujan Bound: A Separation of Cut Versus Spectral Sparsification. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3693–3731. SIAM, 2022.
- [DMS17] Amir Dembo, Andrea Montanari, and Subhabrata Sen. Extremal cuts of sparse random graphs. *The Annals of Probability*, 45(2):1190–1217, 2017.
- [DPMTL21] Giacomo De Palma, Milad Marvian, Dario Trevisan, and Seth Lloyd. The quantum Wasserstein distance of order 1. *IEEE Transactions on Information Theory*, 67(10):6627–6643, 2021.
- [DS04] Luca De Sanctis. Random multi-overlap structures and cavity fields in diluted spin glasses. *Journal of statistical physics*, 117(5):785–799, 2004.
- [DSS16] Jian Ding, Allan Sly, and Nike Sun. Satisfiability threshold for random regular NAE-SAT. *Communications in Mathematical Physics*, 341(2):435–489, 2016.
- [EAMS21] Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Optimization of mean-field spin glasses. *The Annals of Probability*, 49(6):2922–2960, 2021.
- [EWL⁺21] Sepehr Ebadi, Tout T Wang, Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Dolev Bluvstein, Rhine Samajdar, Hannes Pichler, Wen Wei Ho, et al. Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature*, 595(7866):227–232, 2021.
- [FGG14] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.

- [FGG20a] Edward Farhi, David Gamarnik, and Sam Gutmann. The Quantum Approximate Optimization Algorithm Needs to See the Whole Graph: A Typical Case. *arXiv preprint arXiv:2004.09002*, 2020.
- [FGG20b] Edward Farhi, David Gamarnik, and Sam Gutmann. The Quantum Approximate Optimization Algorithm Needs to See the Whole Graph: Worst Case Examples. *arXiv preprint arXiv:2005.08747*, 2020.
- [FGGZ19] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Leo Zhou. The Quantum Approximate Optimization Algorithm and the Sherrington-Kirkpatrick Model at Infinite Size. *arXiv preprint arXiv:1910.08187*, 2019.
- [FGL12] Xiequan Fan, Ion Grama, and Quansheng Liu. Hoeffding’s inequality for supermartingales. *Stochastic Processes and their Applications*, 122(10):3545–3559, 2012.
- [FL03] Silvio Franz and Michele Leone. Replica bounds for optimization problems and diluted spin systems. *Journal of Statistical Physics*, 111(3):535–564, 2003.
- [GJ21] David Gamarnik and Aukosh Jagannath. The overlap gap property and approximate message passing algorithms for p -spin models. *The Annals of Probability*, 49(1):180–205, 2021.
- [GJW20] David Gamarnik, Aukosh Jagannath, and Alexander S Wein. Low-Degree Hardness of Random Optimization Problems. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 131–140. IEEE Computer Society, 2020.
- [GS14] David Gamarnik and Madhu Sudan. Limits of local algorithms over sparse random graphs. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 369–376, 2014.
- [GT02] Francesco Guerra and Fabio Lucio Toninelli. The thermodynamic limit in mean field spin glass models. *Communications in Mathematical Physics*, 230(1):71–79, 2002.
- [GT04] Francesco Guerra and Fabio Lucio Toninelli. The high temperature region of the Viana–Bray diluted spin glass model. *Journal of statistical physics*, 115(1):531–555, 2004.
- [GWZ⁺21] Ming Gong, Shiyu Wang, Chen Zha, Ming-Cheng Chen, He-Liang Huang, Yulin Wu, Qingling Zhu, Youwei Zhao, Shaowei Li, Shaojun Guo, et al. Quantum walks on a programmable two-dimensional 62-qubit superconducting processor. *Science*, 372(6545):948–952, 2021.
- [GZ19] David Gamarnik and Ilias Zadik. The landscape of the planted clique problem: Dense subgraphs and the overlap gap property. *arXiv preprint arXiv:1904.07174*, 2019.
- [H⁺63] Theodore Edward Harris et al. *The theory of branching processes*, volume 6. Springer Berlin, 1963.

- [Has19] Matthew B Hastings. Classical and quantum bounded depth approximation algorithms. *arXiv preprint arXiv:1905.07047*, 2019.
- [KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007.
- [Kum92] Vipin Kumar. Algorithms for constraint-satisfaction problems: A survey. *AI magazine*, 13(1):32–32, 1992.
- [KV05] Subhash Khot and Nisheeth K Vishnoi. On the unique games conjecture. In *FOCS*, volume 5, page 3. Citeseer, 2005.
- [Mar21] Kunal Marwaha. Local classical MAX-CUT algorithm outperforms $p = 2$ QAOA on high-girth regular graphs. *Quantum*, 5:437, 2021.
- [Mon19] A. Montanari. Optimization of the Sherrington-Kirkpatrick Hamiltonian. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1417–1433, Los Alamitos, CA, USA, nov 2019. IEEE Computer Society.
- [MP01] Marc Mézard and Giorgio Parisi. The Bethe lattice spin glass revisited. *The European Physical Journal B-Condensed Matter and Complex Systems*, 20(2):217–233, 2001.
- [Pan14a] Dmitry Panchenko. Introduction to the SK model. *arXiv preprint arXiv:1412.0170*, 2014.
- [Pan14b] Dmitry Panchenko. The Parisi formula for mixed p -spin models. *The Annals of Probability*, 42(3):946–958, 2014.
- [Par80] Giorgio Parisi. A sequence of approximated solutions to the SK model for spin glasses. *Journal of Physics A: Mathematical and General*, 13(4):L115, 1980.
- [Pre18] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.
- [PRTR14] Giorgio Parisi, Federico Ricci-Tersenghi, and Tommaso Rizzo. Diluted mean-field spin-glass models at criticality. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(4):P04013, 2014.
- [PT04] Dmitry Panchenko and Michel Talagrand. Bounds for diluted mean-fields spin glass models. *Probability Theory and Related Fields*, 130(3):319–336, 2004.
- [Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 245–254, 2008.
- [Sen18] Subhabrata Sen. Optimization on sparse random hypergraphs and spin glasses. *Random Structures & Algorithms*, 53(3):504–536, 2018.
- [SK75] David Sherrington and Scott Kirkpatrick. Solvable model of a spin-glass. *Physical review letters*, 35(26):1792, 1975.

- [Tak90] Lajos Takács. On Cayley’s formula for counting forests. *Journal of Combinatorial Theory, Series A*, 53(2):321–323, 1990.
- [Tal06] Michel Talagrand. The parisi formula. *Annals of mathematics*, pages 221–263, 2006.
- [Wei22] Alexander S Wein. Optimal low-degree hardness of maximum independent set. *Mathematical Statistics and Learning*, 2022.
- [WL20] Jonathan Wurtz and Peter J Love. Bounds on MAXCUT QAOA performance for $p > 1$. *arXiv preprint arXiv:2010.11209*, 2020.
- [YFW⁺03] Jonathan S Yedidia, William T Freeman, Yair Weiss, et al. Understanding belief propagation and its generalizations. *Exploring artificial intelligence in the new millennium*, 8:236–239, 2003.
- [ZWC⁺20] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D Lukin. Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices. *Physical Review X*, 10(2):021067, 2020.

A Proof of Proposition 3.2

Proof. We describe a circuit that implements a Bell experiment: a Bell pair of entangled qubits is created, and then a unitary transformation is randomly and independently applied to each qubit before they are measured (equivalently, a random basis is chosen for each measurement).

Consider 4 qubits $|a_c a_e b_e b_c\rangle$ in the state $|0\rangle^{\otimes 4}$. Apply the H gate to a_c , a_e and b_c so that they enter into the $|+\rangle$ state, and leave b_e as is. Using a_e as the control qubit, apply a CNOT gate to b_e . This results in the creation of a Bell pair $|\phi_+\rangle$ between a_e and b_e . Now apply a controlled unitary $C-U$ to a_e using a_c as the control qubit. Similarly, apply the controlled unitary $C-U$ to b_e using b_c as the control qubit. Finally, measure all qubits in the Z basis.

Now, form a graph over the qubits of this circuit, with edges between pairs of qubits that are interacted on by the same gate, as well as a self-loop on b_e so as to distinguish it from a_e :

$$G = (\{a_c, a_e, b_e, b_c\}, \{\{a_c, a_e\}, \{a_e, b_e\}, \{b_e, b_c\}, \{b_e\}\}).$$

An 1-local algorithm on graphs which implements the above quantum circuit when run on G is as follows:

1. Create a qubit for each vertex, in the $|0\rangle$ state.
2. Apply a Hadamard gate to each vertex without a self-loop.
3. For every edge between a vertex of degree 4 with a self-loop and a vertex of degree 2, apply a controlled-not gate from the vertex of degree 2 to the one of degree 4.
4. For every edge incident to a vertex of degree 1, apply a controlled-U gate from that vertex to the one at the other endpoint of the edge.
5. Measure all qubits in the Z basis and output the results.




Figure 2: A graph G (left) on which the algorithm described in the proof of Proposition 3.2 executes a quantum circuit (right) which performs a Bell experiment. a_e and b_e are entangled, as suggested by their names. a_c and b_c are independent and correspond to control qubits.

We show that this algorithm is 1-local by doing casework on each "type" of vertex. It is ok for the output on a vertex to depend on the degree of the vertex and whether it has any self-loops, as these are functions of the 1-neighborhood of the vertex. Thus the types of vertices are as follows:

Degree-1 vertices The output of a vertex of degree 1 is always an independent uniform distribution over $\{-1, 1\}$. The only interaction of the corresponding qubit is being the source of a controlled- U gate. This gate commutes with Pauli- Z operators on the source qubit. Therefore, the measurement in the Z basis commutes past the controlled- U gate, effectively making this part of a circuit a Z measurement on a $|+\rangle$ state, followed by classical control determining whether to apply a U gate on its neighbor. Since this measurement is always independent and invariant with respect to the hypergraph, the algorithm is 1-local on these vertices.

Degree-2 vertices These can interact with degree-1 vertices and degree-4 vertices with self-loops. If we consider the Heisenberg picture and propagate the Z measurements backward through the circuit, the first gates we encounter are controlled- U gates from the degree-1 vertices. Recall that we can treat those controlled- U gates as a single-qubit U gate controlled classically by a single random bit. If the U gate happens, then a Z measurement after the U gate is equivalent to a $U^\dagger Z U$ measurement before the U gate; otherwise it remains a Z measurement. If there's more than one degree-1 neighbor and more than one consequent U gate, then our measurement is $U^{2^\dagger} Z U^2$ instead, and we've reached the initialization of our qubit. If the other neighbor is a degree-4 vertex with a self-edge, then there's a controlled-NOT acting on our degree-2 vertex, and $U^\dagger Z U$ after the controlled-NOT is equivalent to some linear combination of phase changes on the degree-4 vertex multiplied by some single-qubit unitary on the degree-2 vertex. At this point, our measurement operator has reached the initializations of 2 qubits, both in the 1-neighborhood of the degree-2 vertex we started with. If there was no U gate applied in step 4, and only one or more controlled-NOTs instead, then a Z on the target after a controlled-NOT gate is equivalent to the product of Z on the target and Z on the source before the gate. Then the Z on the target will reach an initialization, while the Z on the source commutes past other being on the source of other controlled-NOT gates, to also reach the initialization. Thus we have described how the measurement is a function of edges in the 1-neighborhood, and how the support of the measurement in the Heisenberg picture is within the 1-neighborhood, so that if all other measurements also have supports in their respective 1-neighborhoods, the algorithm is 1-local on vertices of degree 2.

Degree-4 vertices These can interact with degree-1 vertices and degree-2 vertices, but only if they have a self-loop. The action of degree-1 vertices (in step 4 of the algorithm) is identical with the degree-2 case above. If the backwards-propagating measurement is Z at the time when it hits the source of a controlled-NOT gate, the Z on the source after the controlled-NOT is equivalent to a Z on the source multiplied by an X on the target before the gate. The Z on the source then propagates back through another controlled-NOT if it exists, remaining a Z until it hits the initialization. The X s on the target degree-2 vertices commute through at most one other controlled-NOT gate on the target end until it hits the initialization of the degree-2 vertex. At the end of this process, the support of the measurement operator on the qubit initializations is fully determined by and restricted to the 1-neighborhood of the vertex. On the other hand, if the measurement is $U^\dagger Z U$ by the time it touches the source of a controlled-not, it propagates backwards into a linear combination of single-qubit unitaries on the source

and I or X operators on the target. The unitary on the source then hits the initialization of the degree-4 vertex, while the I or the X commutes past any other controlled-NOTs the degree-2 vertex might be the target of, to hit the initialization of the adjacent degree-2 vertex. Again, the measurements in the Heisenberg picture are fully determined by and restricted to the 1-neighborhood.

All other degrees Vertices of any other degree do not interact with other vertices in the algorithm.

By the above casework, this algorithm is 1-local.

Note that this process *cannot* be encoded in a 1-local factors of i.i.d. algorithm since this setup allows for signaling strategies that violate Bell inequalities, whereas 1-local factors of i.i.d. can be explained by using latent variables to describe the evolution of the randomness of the 1-neighborhood of every vertex. An extension of this circuit which would involve generating p entangled qubits in a similar process generalizes the argument to p -local algorithms. $\square$

B Proof of Theorem 2.20

Proof. First, note that given $\frac{dn}{k}$ hyperedges (which is the expectation of the $\text{Poisson}(dn/k)$ distribution from which the number of edges are sampled), each of size k , the expected number of hyperedges some vertex v_i shows up in is,

$$\mathbb{E}_{|E|} [\Pi_i(E(HG))] = \frac{dn}{k} \cdot \Pr_{e \sim [n]^k} [v_i \in e] \leq \frac{dn}{k} \cdot \left(\frac{n^{k-1}}{n^k} \right) k = \frac{dn}{k} \cdot \left(\frac{1}{n} \right) k = d. \quad (\text{B.1})$$

Now, consider another model of a k -uniform hypergraph in which we sample n^k edges independently, each with probability p . This induces a $\text{Bin}(n^k, p)$ distribution on the number of hyperedges. To compare this with our model, we compare the expected number of hyperedges as,

$$pn^k = \frac{dn}{k} \implies p = \frac{d}{kn^{k-1}}.$$

Note that the degree distribution of a vertex (which is equivalent to the number of hyperedges it appears in) in this model is given as,

$$\text{Bin} \left(kn^{k-1}, \frac{d}{kn^{k-1}} \right).$$

The distribution above converges to $\text{Poisson}(d)$ in the large n limit, and its moment generating function is dominated by that of the $\text{Poisson}(d)$ distribution for all large but finite n . Formally,

$$\phi \left(\text{Bin} \left(kn^{k-1}, \frac{d}{kn^{k-1}} \right) \right) = (1 - d(1 - e^t))^{kn^{k-1}} \leq e^{d(e^t - 1)} = \phi(\text{Poisson}(d)).$$

Therefore, the number of vertices in the 1-neighborhood of any vertex $v \in V(G)$ can be bounded from above (in the large n limit) as,

$$(k-1)\text{Poisson}(d).$$

So, we will consider the scaled Galton-Watson process above starting at some vertex $v \in V(G)$, which is itself a Galton-Watson process [H⁺63]. We upper bound the scaled Galton-Watson process above with the Galton-Watson process induced by the $\text{Poisson}(d(k-1))$ distribution (at an appropriate level) by comparing their respective probability generating functions. Note that the probability generating function of $(k-1)\text{Poisson}(d)$ is $f_{\text{Poisson}(d)}(z^{k-1})$. To look at the x -th neighborhood, we will use the probability generating function of the x -th level of the underlying Galton-Watson process, denoted as f_x . Now, by [H⁺63, Theorem 4.1],

$$f_x = f_1^{\circ x} := \overbrace{f_1 \circ \cdots \circ f_1}^x,$$

where $f_1 = f_{\text{Poisson}(d)}(z^{k-1}) = f_{(k-1)\text{Poisson}(d)}(z)$. By comparing the pgf of the Galton-Watson process of $f_{(k-1)\text{Poisson}(d)}(z)$ and the Galton-Watson process of $f_{\text{Poisson}(d(k-1))}(z)$, one can observe that,

$$f_{x-1} \leq g_x,$$

where $g_x = g_1^{\circ x}$, and $g_1 = f_{\text{Poisson}(d(k-1))}(z)$. We now repeat the argument in [FGG20a, Neighborhood Theorem] that bounds the size of a $\text{Poisson}(d(k-1))$ branching process.

Let Z_x denote the size of the x -th generation of a $\text{Poisson}(d(k-1))$ branching process. Note that $Z_0 = 1$, $Z_1 = \text{Poisson}(d(k-1))$ and, more generally, one can look at the moment generating function of the $\text{Poisson}(d(k-1))$ branching process as,

$$\mathbb{E}_{\text{Branching}(d(k-1))} [e^{tZ_x}] = e^{d(k-1)(\mathbb{E}[e^{tZ_{x-1}}] - 1)}.$$

We denote by $\phi_x(t)$ the moment generating function of the $\text{Poisson}(d(k-1))$ branching process. It is straightforward to see by an inductive argument used in [FGG20a, Neighborhood Size Theorem] that,

$$\phi_x \left(\left(\frac{\ln 2}{d(k-1)} \right)^x \right) \leq e, \quad \forall k \geq 0.$$

Furthermore, by an application of Markov's inequality to the moment generating function, the following is true for any $u, t > 0$,

$$\Pr_{\text{Branching}(d(k-1))} \left[Z_x \geq u \left(\frac{(d(k-1))}{\ln 2} \right)^x \right] \leq e^{-u \left(\frac{d(k-1)}{\ln 2} \right)^x} \phi_x(t) \leq e^{-u} e, \quad (\text{B.2})$$

where the last inequality follows by a choice of $t = \left(\frac{\ln 2}{d(k-1)} \right)^x$.

We now bound the probability that the total number of nodes in the branching process at height x is *at least* c via a union bound,

$$\Pr_{\text{Branching}(d(k-1))} \left[\sum_{i=1}^x Z_i \geq c \right] \leq \sum_{i=1}^x \Pr \left[Z_i \geq \frac{c}{x} \right].$$

Let $c = (d(k-1))^{sx} \left(\frac{d(k-1)}{\ln 2} \right)^x$ and $u = \frac{d^{sx}}{x}$, where s will be chosen later to demonstrate the existence of A . Substituting these into Equation B.2 yields a lower bound on the size of the neighborhoods induced by the $\text{Poisson}(d(k-1))$ branching process as,

$$\Pr_{\text{Branching}(d(k-1))} \left[\sum_{i=1}^x Z_i \geq (d(k-1))^{sx} \left(\frac{d(k-1)}{\ln 2} \right)^x \right] \leq x e^{-\frac{d^{sx}}{x}} e \leq e^{-d^{sx}/2}, \quad (\text{B.3})$$

where we assume a sufficiently large choice of x . Denote by $p_r(v_i)$ the probability that a vertex v_i in a random k -uniform hypergraph with $p = \frac{d}{kn^{k-1}}$ has a x -neighborhood with size *at least* that of Equation B.3, conditioned on the hypergraph having r hyperedges. Extending the argument in [FGG20a, Neighborhood Size Theorem] further,

$$\sum_{r=dn/k}^{\infty} \Pr[\text{G has } r \text{ hyperedges}] \cdot p_r(v_i) \leq \sum_{r=0}^{\infty} \Pr[\text{G has } r \text{ hyperedges}] \cdot p_r(v_i) \leq e^{-d^{sx}/2}.$$

Since an increase in the number of sampled edges will only increase $p_r(v_i)$, it follows that,

$$\sum_{r=dn/k}^{\infty} \Pr[\text{G has } r \text{ hyperedges}] \cdot p_{dn/k}(v_i) \leq e^{-d^{sx}/2}.$$

Now, the number of edges are distributed as $\text{Poisson}(dn/k)$. Since $\mathbb{E}[\text{Poisson}(dn/k)] = \frac{dn}{k}$,

$$\sum_{r=dn/k}^{\infty} \Pr[\text{G has } r \text{ hyperedges}] = \Pr[\text{G has } \geq dn/k \text{ hyperedges}].$$

Note that the mean of the $\text{Poisson}(dn/k)$ distribution is an integer. Consequently, the median of the distribution is also $\frac{dn}{k}$ [Cho94]. Then,

$$\Pr[\text{Poisson}(dn/k) \geq dn/k] \in \left(\frac{1}{2} - \Pr[\text{Poisson}(dn/k) = dn/k], \frac{1}{2} + \Pr[\text{Poisson}(dn/k) = dn/k] \right).$$

Now, by applying Stirling's approximation,

$$\Pr[\text{Poisson}(dn/k) = dn/k] \leq \frac{1}{\sqrt{2\pi dn/k}}.$$

Consequently,

$$\Pr[\text{Poisson}(dn/k) \geq dn/k] \geq \frac{1}{2} - \Pr[\text{Poisson}(dn/k) = dn/k] \geq \frac{1}{2} - \frac{1}{\sqrt{2\pi dn/k}} = \frac{1}{2} - o_{d,n}(1).$$

This yields an upper bound for the probability of a x -neighborhood exceeding the desired size as,

$$p_{dn/k}(v_i) \leq (2 + o_{d,n}(1)) e^{-d^{sx}/2} \leq e^{-d^{sx}/3},$$

for sufficiently large x .

We set

$$2p \leq x - 1 = \frac{(1 - \tau) \log n}{\log \left(\frac{d(k-1)}{\ln 2} \right)} - 1$$

for some $\tau \in (0, 1)$. Then, let $\log = \log_{d(k-1)}$ and define $L = \log_{d(k-1)} \left(\frac{1}{\ln 2} \right)$. Consequently, the remaining argument follows exactly as in the last part of the proof of [FGG20a, Neighborhood Size Theorem]. Specifically, Equation B.3 reduces to [FGG20a, Eq. 75, Neighborhood Size Theorem] with $w = (1 - \tau)$ given the parameter choices above after some algebra. $\square$