

The Algorithmic Phase Transition of Random k -SAT for Low Degree Polynomials

Guy Bresler*

Brice Huang†

June 3, 2021

Abstract

Let Φ be a uniformly random k -SAT formula with n variables and m clauses. We study the algorithmic task of finding a satisfying assignment of Φ . It is known that a satisfying assignment exists with high probability at clause density $m/n < 2^k \log 2 - \frac{1}{2}(\log 2 + 1) + o_k(1)$, while the best polynomial-time algorithm known, the **Fix** algorithm of Coja-Oghlan [25], finds a satisfying assignment at the much lower clause density $(1 - o_k(1))2^k \log k/k$. This prompts the question: is it possible to efficiently find a satisfying assignment at higher clause densities?

To understand the algorithmic threshold of random k -SAT, we study low degree polynomial algorithms, which are a powerful class of algorithms including **Fix**, Survey Propagation guided decimation (with bounded or mildly growing number of message passing rounds), and paradigms such as message passing and local graph algorithms. We show that low degree polynomial algorithms can find a satisfying assignment at clause density $(1 - o_k(1))2^k \log k/k$, matching **Fix**, and not at clause density $(1 + o_k(1))\kappa^* 2^k \log k/k$, where $\kappa^* \approx 4.911$. This shows the first sharp (up to constant factor) computational phase transition of random k -SAT for a class of algorithms. Our proof establishes and leverages a new many-way overlap gap property tailored to random k -SAT.

1 Introduction

The k -SAT problem occupies a central role in complexity theory as the first and canonical NP-complete problem [30]. Its average-case analogue, random k -SAT, likewise has an important role in average-case computational complexity. In this problem, we are given a k -SAT formula with m clauses and n variables, whose km literals are each sampled i.i.d. from the set of $2n$ possible literals, and our task is to produce a satisfying assignment [1].¹ There are two natural basic questions for random k -SAT. First, at what scalings of (n, k, m) are there satisfying assignments? Second, when can they be found by efficient algorithms?

Early work showed that the interesting regime of random k -SAT (for fixed k) is when $m = \Theta(n)$, and that the problem's qualitative behavior in the large- n limit depends on the *clause density* (or *clause-to-variable ratio*) $\alpha = m/n$. Specifically, [37] showed that if $\alpha \geq 2^k \log 2$, the random k -SAT formula is unsatisfiable with high probability, while [18] showed that if $\alpha < 2^k/k$, a simple algorithm (unit clause propagation) will find a satisfying assignment with nontrivial probability. As we tune α , we encounter phase transitions separating one qualitative behavior from another. Two such phase transitions are of primary interest to us: the *satisfiability threshold*, below which the k -SAT formula admits a satisfying assignment with high probability, and the *algorithmic threshold*, below which there exists a polynomial-time algorithm producing a satisfying assignment with high probability. The satisfiability threshold is well understood: Kirousis,

*Massachusetts Institute of Technology, Department of EECS. Email: guy@mit.edu. Supported by MIT-IBM Watson AI Lab and NSF CAREER award CCF-1940205.

†Massachusetts Institute of Technology, Department of EECS. Email: bmhuang@mit.edu. Supported by NSF Graduate Research Scholarship 1745302, a Siebel Scholarship, and NSF TRIPODS award 1740751.

Key words and phrases. Random k -SAT; average-case complexity; algorithmic phase transition; low degree polynomial hardness; overlap gap property.

¹In a variant of this definition, the m clauses are chosen uniformly and without replacement among all $2^k \binom{n}{k}$ clauses with k distinct, non-complementary literals. This definition behaves identically to ours in the large- n limit, and all properties of random k -SAT we show in this paper apply equally to this model.

Kranakis, Krizanc, and Stamatiou [58] showed that with high probability a satisfying assignment does not exist at clause density $2^k \log 2 - \frac{1}{2}(\log 2 + 1) + o_k(1)$, where $o_k(1)$ denotes a term limiting to 0 as $k \rightarrow \infty$, while Coja-Oghlan and Panagiotou [29] showed that a satisfying assignment exists with high probability at clause density $2^k \log 2 - \frac{1}{2}(\log 2 + 1) - o_k(1)$. Ding, Sly, and Sun [33] closed this $o_k(1)$ gap, showing that for all sufficiently large k random k -SAT sharply transitions from satisfiable to unsatisfiable at a threshold clause density α_s which they determined.

In contrast, the algorithmic threshold is poorly understood. The current best polynomial-time algorithm, the **Fix** algorithm of Coja-Oghlan [25], can find a satisfying assignment with high probability at clause density $(1 - o_k(1))2^k \log k/k$, nearly a factor of k below the satisfiability threshold. A body of evidence, both rigorous and non-rigorous, has emerged to suggest that this is the correct threshold, but until now there has been no result ruling out any class of algorithms at this threshold.

In the early 2000s, statistical physicists developed a rich but non-rigorous theory describing the solution geometry of random k -SAT, among other random constraint satisfaction problems [60]. This theory predicts several phase transitions in the solution geometry as the clause density increases. At low clause density, the random k -SAT instance is satisfiable by many assignments, which form one large cluster. When the clause density increases past the *uniqueness threshold*, some additional disconnected solution clusters appear, but the large cluster contains all but an exponentially small fraction of solutions. Past the *clustering threshold*, the solution space shatters into an exponentially large number of disconnected clusters, each with an exponentially small fraction of solutions. Past the *condensation threshold*, while there remain an exponential number of clusters, the mass of the solutions is dominated by a few large clusters. Finally, beyond the *satisfiability threshold* there are no satisfying assignments with high probability. For a pictorial description of these phase transitions, see [60, Figure 2]. The authors of [60] also predicted that Markov Chain Monte Carlo (MCMC) algorithms succeed up to the clustering threshold and no more. Their prediction of the satisfiability threshold was later confirmed by the result of Ding, Sly and Sun [33] mentioned above. Moreover, the physics prediction of the condensation threshold for random regular NAE- k -SAT was recently confirmed by Nam, Sly, and Sohn [66].

Since then, the clustering threshold, which is predicted to be (suppressing lower order terms; see [60, Equation 6] for a more precise expression) $(1 + o_k(1))2^k \log k/k$, has emerged as the predicted limit of *all* efficient algorithms. Above this threshold, structural phenomena of the solution space have been rigorously established that (still non-rigorously) suggest algorithmic hardness. Achlioptas and Coja-Oghlan [3] showed that above clause density $(1 + o_k(1))2^k \log k/k$, for some setting of the $o_k(1)$ term, long-range correlations appear in a k -SAT instance's solution space, in the following sense. Say variable x_i of a satisfying assignment $x \in \{\mathbf{T}, \mathbf{F}\}^n$ is *frozen* if any satisfying assignment y with $x_i \neq y_i$ is at Hamming distance $\Omega(n)$ from x . Above this clause density, in all but an $o(1)$ fraction of satisfying assignments, all but an $o_k(1)$ fraction of bits are frozen with high probability. This suggests that above clause density $(1 + o_k(1))2^k \log k/k$, local search algorithms are unlikely to succeed, and any algorithmic solution to random k -SAT must use a qualitatively different approach.

The rigorous evidence for the algorithmic threshold consists of exhibiting algorithms on one side and producing bounds against specific algorithms or restricted computational models on the other side. There is a long line of work on heuristic algorithms for k -SAT. The oldest such heuristic is the *Davis-Putnam-Logemann-Loveland* (DPLL) algorithm [32, 31], a backtracking based search algorithm which still forms the basis for many modern SAT solvers. Other heuristics that have emerged are the *pure literal rule* [48]; *unit clause propagation* [18]; *shortest clause* [23, 39]; *walksat* [68, 27]; and *Belief Propagation* and *Survey Propagation guided decimation* [65, 14]. However, there is no evidence, rigorous or non-rigorous, that any of these algorithms succeed with high probability beyond clause density $O_k(2^k/k)$. The breakthrough result of Coja-Oghlan [25] produced the algorithm **Fix**, which is proven to succeed up to clause density $(1 - o_k(1))2^k \log k/k$ with high probability. This is the best algorithm to date, and the above evidence from physics suggests that this clause density is optimal.

In the way of rigorous negative results, the earliest work is by Luby, Mitzenmacher, and Shokrollahi [62], who prove that the pure literal rule does not solve random 3-SAT above clause density approximately 1.63. Achlioptas and Sorkin [4] generalized this result, showing that the class of so-called *myopic algorithms* cannot solve random 3-SAT above clause density approximately 3.26. For large k , the earliest work is by Achlioptas, Beame, and Molloy [2], who show that beyond clause density $O_k(2^k/k)$, DPLL type algorithms require an exponential running time. Gamarnik and Sudan [46] showed that the class of *balanced sequential*

local algorithms, which includes Belief Propagation and Survey Propagation guided decimation (with constant or mildly growing number of message passing rounds) fail to solve random NAE- k -SAT at clause density $(1 + o_k(1))2^{k-1} \log^2 k/k$. The quantity 2^{k-1} is the NAE- k -SAT analogue of 2^k for k -SAT. The remaining negative results in the literature are bounds against specific algorithms, proved by tailored analysis. Hetterich [52] proved that Survey Propagation guided decimation (without restriction on the number of message passing rounds) fails at clause density $(1 + o_k(1))2^k \log k/k$, and Coja-Oghlan, Haqshenas, and Hetterich [28] proved that walksat fails at clause density $O_k(2^k \log^2 k/k)$. To date, all negative results in the literature have been either not asymptotically tight against the conjectured threshold $(1 + o_k(1))2^k \log k/k$ or tailored to a particular algorithm.

In this paper, we show the first hardness result for a class of algorithms that is asymptotically tight against the conjectured threshold. We show that the class of low degree polynomial algorithms do not solve random k -SAT past clause density $(1 + o_k(1))\kappa^* 2^k \log k/k$, for a constant $\kappa^* \approx 4.911$. Low degree polynomial algorithms include many of the above algorithms, including **Fix**, sequential local algorithms, message passing algorithms, and local algorithms on the factor graph. By confirming the physics view to a large extent, this gives strong evidence that $(1 + o_k(1))2^k \log k/k$ is the correct algorithmic threshold.

While our hardness result is meaningful in its own right, the machinery we use to prove it is significant as well. Our proof is based on making rigorous an appropriate understanding of random k -SAT's energy landscape. This proof extends a line of work on the *overlap gap property* (OGP) and develops techniques to overcome obstacles limiting the reach of prior OGP methodology. We now briefly summarize the OGP program and our contribution to it; a more detailed discussion can be found in Section 3.

The OGP program, initiated by Gamarnik and Sudan in [45], is the first line of work to translate physics intuitions about solution geometry to rigorous results ruling out classes of algorithms. In its basic form, an OGP argument shows that beyond some phase transition, with high probability any two solutions to an average-case problem have either small or large overlap; the exclusion of medium overlaps is one formalization of the clustering phenomenon. The argument then shows that a smooth algorithm solving the problem can be used to construct the forbidden structure, and thus such an algorithm cannot exist. To improve the threshold at which algorithms are ruled out, these arguments have been generalized to consider forbidden structures consisting of several solutions, which we term multi-OGPs. Using multi-OGPs, a line of work [69, 43] culminating in the paper of Wein [71] tightly identified the algorithmic phase transition of maximum independent set on a sparse Erdős-Rényi graph for low degree polynomials. Multi-OGPs have also been used in [46, 44] to rule out classes of algorithms for random k -SAT and the number partitioning problem well below the point where solutions exist, although these results are not tight against the best algorithms.

Establishing a multi-OGP for random k -SAT that produces a tight hardness result is far more challenging than for maximum independent set. The difficulty lies in analyzing the free energy of an overlap structure consisting of several satisfying assignments; to establish a multi-OGP, we must find an overlap structure making this free energy negative. The solution geometry of maximum independent set made the analogous free energy relatively simple, which made that problem amenable to OGP. In contrast, the free energy for random k -SAT has complex dependencies which make this analysis difficult, and it is a priori not even clear how to define the forbidden structure for which the multi-OGP holds. We identify the correct forbidden structure and prove that with high probability it does not occur. To achieve this, we make three conceptual contributions. First, we define a notion of overlap profile and conditional overlap entropy. Second, we define the multi-OGP in terms of this formalism; this is itself a key innovation, as all the (multi-)OGPs in the literature have not required the full power of the overlap profile, and the forbidden structure we use is more intricate than those in the literature. Third, we perform a novel free energy analysis to show the desired multi-OGP occurs. We are optimistic that many more problems, including those with complicated free energies, may be amenable to the techniques developed in this paper.

Reasoning about the power of restricted classes of algorithms is at the heart of theoretical computer science. As discussed above, there is a line of work showing hardness (at suboptimal clause densities) of random k -SAT against restricted computational models [2, 46]. More generally, for other problems, the limits of various other restricted classes of algorithms have been studied, including circuits [5, 40, 50, 20], the convex hierarchies of Sherali-Adams and Lóvász-Schrijver (see [19] and references therein), the sum of squares hierarchy [49, 59, 9], and local algorithms on graphs [45, 69]. Recently, the class of low degree polynomial algorithms has emerged as a prominent class in average case complexity and statistical inference. As sketched in [43, Appendix A], this class contains many popular and powerful frameworks, including local

algorithms on graphs, power iteration, and approximate message passing [35, 11, 56, 64, 36]. In addition, a recent flurry of work has shown that for many average-case problems in high-dimensional statistics, including planted clique, sparse PCA, community detection, and tensor PCA, low degree polynomials are as powerful as the best polynomial-time algorithms known [55, 54, 53, 7, 61, 34, 22, 15, 63, 70, 8, 16]. Thus, showing that low degree polynomial algorithms fail at some threshold provides evidence that all polynomial-time algorithms fail at that threshold.

1.1 Notation

For all positive integers n , $[n]$ denotes the set $\{1, \dots, n\}$. For two assignments $x, y \in \{\mathbf{T}, \mathbf{F}\}^n$, let $\Delta(x, y) = \frac{1}{n} |\{i \in [n] : x_i \neq y_i\}|$ denote the normalized Hamming distance between x and y .

Throughout, $\log$ denotes the natural logarithm. The binary entropy function $H : [0, 1] \rightarrow [0, 1]$ is given by $H(x) = -x \log x - (1 - x) \log(1 - x)$. We will often use the basic inequality $H(x) \leq x \log \frac{e}{x}$. We will also overload notation and use $H(\cdot)$ to denote the entropy of certain distributions, for instance $H(\pi)$. These will be carefully defined where first used.

All our results are in the double limit as $n \rightarrow \infty$, and then $k \rightarrow \infty$. Thus, the notations $O(\cdot)$, $\Omega(\cdot)$, $o(\cdot)$, $\omega(\cdot)$ indicate asymptotic behavior in n , suppressing any dependence on k . When subscripted with k , these notations indicate asymptotic behavior in k of a quantity independent of n .

Organization. The rest of this paper is structured as follows. In Section 2 we state our main impossibility and achievability results, Theorems 2.6 and 2.10, that low degree polynomial algorithms cannot solve random k -SAT beyond clause density $(1 + o_k(1))\kappa^* 2^k \log k/k$ and can solve random k -SAT at clause density $(1 - o_k(1))2^k \log k/k$. In Section 3 we summarize the progress of the OGP program and place our work's contributions in context. Sections 4 through 6 are devoted to the proof of Theorem 2.6. Section 4 develops the formalism needed to define our central multi-OGP and outlines the proof of Theorem 2.6. It proves this theorem assuming Propositions 4.7(a) and 4.7(c), which control the probabilities of the low degree polynomial output being stable and of the multi-OGP occurring. Section 5 proves Proposition 4.7(c), showing that the multi-OGP occurs. The proofs in this section contain many of our main technical contributions. Section 6 proves Proposition 4.7(a), completing the proof of Theorem 2.6. Finally, Section 7 shows our converse achievability result, Theorem 2.10, by showing that low degree polynomials can simulate Fix .

Acknowledgements. We are grateful to David Gamarnik for useful conversations. BH is also grateful to Alex Wein and Mehtaab Sawhney for useful discussions. This work was done in part while the authors were participating in the Probability, Geometry, and Computation in High Dimensions program at the Simons Institute for the Theory of Computing in Fall 2020.

2 Results

Throughout this paper, let $\mathcal{V} = \{x_1, \dots, x_n\}$ denote a set of propositional variables. The set of corresponding literals, consisting of the variables in $\mathcal{V}$ and their negations, is denoted $\mathcal{L} = \{x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n\}$. Let $\Omega_k(n, m)$ be the set of all k -SAT formulas, consisting of an AND of m clauses, each of which is an OR of k literals from $\mathcal{L}$. We allow variables to appear multiple times in the same clause, and for clauses to appear multiple times in the formula. We treat each such formula as an ordered m -tuple of clauses, each of which is an ordered k -tuple of literals. For $\Phi \in \Omega_k(n, m)$, let Φ_i ($1 \leq i \leq m$) denote the i th clause of Φ , and let $\Phi_{i,j}$ ($1 \leq j \leq k$) denote the j th literal of Φ_i . The central object of this paper is the random k -SAT formula, defined as follows.

Definition 2.1 (Random k -SAT). The *random k -SAT* distribution $\Phi_k(n, m)$ is the law of a uniformly random sample from $\Omega_k(n, m)$. Equivalently, we can sample a formula $\Phi \sim \Phi_k(n, m)$ by sampling the literals $\Phi_{i,j}$ ($1 \leq i \leq m$, $1 \leq j \leq k$) i.i.d. from $\text{unif}(\mathcal{L})$.

We will study the class of low degree polynomial algorithms, defined as follows. This is the same computational model considered in [43, 71].

Definition 2.2 (Low degree polynomial). A function $f : \mathbb{R}^N \rightarrow \mathbb{R}^n$ is a *polynomial of degree at most D* if it is of the form

$$f(x) = (f_1(x), \dots, f_n(x)),$$

where each $f_i : \mathbb{R}^N \rightarrow \mathbb{R}$ is a multivariate polynomial (in the ordinary sense) with real coefficients of degree at most D . We will henceforth use “degree- D polynomial” to mean polynomial of degree at most D . A *random degree- D polynomial* $f : \mathbb{R}^N \rightarrow \mathbb{R}^n$ is defined in the same way, except the coefficients are allowed to be random (but independent of the input x). Formally, for some probability space $(\Omega, \mathbb{P}_\omega)$, f is a function $f : \mathbb{R}^N \times \Omega \rightarrow \mathbb{R}^n$ such that for each $\omega \in \Omega$, $f(\cdot, \omega)$ is a degree- D polynomial.

Remark 2.3. We will see in Lemma 4.1 below that randomness does not increase the power of the class of low degree polynomials.

Let us define how to encode a k -SAT formula as an input to a low degree polynomial. Define an arbitrary total order on $\mathcal{L}$. Let $N = m \cdot k \cdot 2n$. We can encode each formula $\Phi \in \Omega_k(n, m)$ as a vector in $\{0, 1\}^N$ by writing it as a collection of indicators $\Phi_{i,j,s}$ for $1 \leq i \leq m$, $1 \leq j \leq k$, and $1 \leq s \leq 2n$, where $\Phi_{i,j,s}$ is the indicator that the literal $\Phi_{i,j}$ is the s th element of $\mathcal{L}$. This encoding is an input to a low degree polynomial. Slightly abusing notation, we will identify Φ with its vector encoding. Note that if $\Phi \sim \Phi_k(n, m)$, then each of the sub-vectors $(\Phi_{i,j,s})_{s=1}^{2n}$ for fixed i, j is drawn uniformly from the set of standard basis vectors in $2n$ dimensions.

Next, let us define how to interpret the output of a polynomial as a Boolean assignment. Informally, outputs that are at least 1 indicate the corresponding bit of the assignment is true, and outputs that are at most -1 indicate the corresponding bit is false. It is important to exclude the interval $(-1, 1)$, so that a small change in the output of the polynomial cannot induce a large change in the interpreted assignment. We allow the algorithm to make mistakes in up to an η fraction of positions (which include all the positions where the polynomial outputs a value in $(-1, 1)$, and possibly others), and we say the algorithm succeeds if there is any way to fill in the mistakes to yield a satisfying assignment. Formally, this rounding procedure is defined as follows.

Definition 2.4 (η -assisted rounding). Let $f : \mathbb{R}^N \times \Omega \rightarrow \mathbb{R}^n$ be a random polynomial, and let $\eta \in (0, 1)$. Let $\mathcal{B}$ be a (computationally unbounded) deterministic subroutine, taking as input $(\Phi, y) \in \mathbb{R}^N \times \mathbb{R}^n$ and outputting (B, z) with $B \subseteq [n]$ and $z \in \{\mathbf{T}, \mathbf{F}\}^B$. We say $\mathcal{A} : \mathbb{R}^N \times \Omega \rightarrow \{\mathbf{T}, \mathbf{F}\}^n$ is an η -assisted rounding of f if there exists $\mathcal{B}$ such that $\mathcal{A}$ runs as follows.

- (1) Set $y = f(\Phi, \omega)$. Set $A = \{i \in [n] : |y_i| < 1\}$ and $(B, z) = \mathcal{B}(\Phi, y)$.
- (2) Check that $B \supset A$ and $|B| \leq \eta n$; otherwise output “fail.”
- (3) Output $x \in \{\mathbf{T}, \mathbf{F}\}^n$, where

$$x_i = \begin{cases} \mathbf{T} & i \notin B \text{ and } y_i \geq 1, \\ \mathbf{F} & i \notin B \text{ and } y_i \leq -1, \\ z_i & i \in B. \end{cases}$$

Note that allowing a computationally unbounded assistant to repair an η fraction of mistakes makes the algorithm class more powerful, which makes an impossibility result against this algorithm class *stronger*. We next define what it means for a low degree polynomial to solve random k -SAT.

Definition 2.5 ((δ, γ, η) -solve). Let $f : \mathbb{R}^N \times \Omega \rightarrow \mathbb{R}^n$ be a random polynomial, and let $\delta, \eta \in (0, 1)$ and $\gamma \geq 1$. We say f (δ, γ, η) -solves $\Phi_k(n, m)$ if there exists an η -assisted rounding $\mathcal{A}$ of f such that the following conditions hold over independent $\Phi \sim \Phi_k(n, m)$ and $\omega \sim (\Omega, \mathbb{P}_\omega)$.

- (a) $\mathbb{P}_{\Phi, \omega}[\mathcal{A}(\Phi, \omega) \text{ satisfies } \Phi] \geq 1 - \delta$.
- (b) $\mathbb{E}_{\Phi, \omega}[\|f(\Phi, \omega)\|_2^2] \leq \gamma n$.

Here, δ is the algorithm’s failure probability. Moreover, γ is a normalization parameter, which we think of as a large constant; this is necessary because without it, the condition that valid outputs of f are outside the interval $(-1, 1)$ becomes meaningless because we can simply scale f by a large constant. An analogous

normalization condition appears in other hardness results against low degree polynomials in the literature [43, 71]. The error tolerance of our rounding scheme is a small constant η , as discussed above.

We define the function $\iota : (1, +\infty) \rightarrow \mathbb{R}$ by

$$\iota(\beta) = \frac{\beta}{1 - \beta e^{-(\beta-1)}}.$$

One can easily check that ι is concave and has minimum value $\kappa^* \approx 4.911$, which is attained at $\beta^* \approx 3.513$, the unique solution to $\beta^2 e^{-(\beta-1)} = 1$ in $(1, +\infty)$.

The following theorem is our main result, showing that no low degree polynomial can solve random k -SAT at clause density $\kappa 2^k \log k/k$ for any $\kappa > \kappa^*$.

Theorem 2.6. *Let $\kappa > \kappa^*$. Let $\alpha = \kappa 2^k \log k/k$ and $m = \lfloor \alpha n \rfloor$. There exists $k^* = k^*(\kappa) > 0$ such that for any $k \geq k^*$, there exists $n^* > 0$, $\eta = \Omega_k(k^{-1})$, $C_1 > 0$, and $C_2 > 0$ (depending on κ, k) such that the following holds. If $n \geq n^*$, $\gamma \geq 1$, $1 \leq D \leq \frac{C_1 n}{\gamma \log n}$ and*

$$\delta \leq \exp(-C_2 \gamma D \log n),$$

then there is no random degree- D polynomial that (δ, γ, η) -solves $\Phi_k(n, m)$.

Several remarks about this result are in order.

Remark 2.7. The only property of low degree polynomials we use is their smoothness, in the sense of Proposition 6.2. Thus Theorem 2.6 rules out any algorithm class satisfying the conclusion of this proposition.

Remark 2.8. The constant κ^* in Theorem 2.6 can likely be optimized further, and improving this constant would tighten the constant-factor gap between this theorem and Theorem 2.10 below. However, without using additional properties of the interpolation path, our methods cannot improve κ^* beyond a constant bounded away from 1, approximately 1.716. Thus, further ideas are needed to close this gap. See Appendix B for a discussion of these points.

In spite of Remark 2.8, we believe that the algorithmic phase transition for low degree polynomials does occur at $(1 + o_k(1))2^k \log k/k$, matching the physics prediction. This is formalized in the following conjecture, which we leave as an open problem.

Conjecture 2.9. *Theorem 2.6 holds for all $\kappa > 1$.*

The following result provides a converse to Theorem 2.6 at clause density $(1 - \varepsilon)2^k \log k/k$ for any $\varepsilon > 0$, giving a lower bound on the algorithmic phase transition within a constant factor of Theorem 2.6 and tight against Conjecture 2.9.

Theorem 2.10. *Let $\varepsilon > 0$. Let $\alpha = (1 - \varepsilon)2^k \log k/k$ and $m = \lfloor \alpha n \rfloor$. There exists $k^* = k^*(\varepsilon) > 0$ such that for any $k \geq k^*$ and any $\eta > k^{-12}$, there exists $n^* > 0$, $D > 0$, $\gamma \geq 1$, and function $\delta : \mathbb{N} \rightarrow [0, 1]$ with $\delta(n) = o(1)$ (depending on ε, k, η) such that the following holds for all $n \geq n^*$. If $\delta = \delta(n)$, there exists a (deterministic) degree- D polynomial that (δ, γ, η) -solves $\Phi_k(n, m)$.*

Remark 2.11. The failure probability δ in Theorem 2.10 is not within the range ruled out by Theorem 2.6. This is likely an artifact of our methods. We prove Theorem 2.10 by simulating (a part of) **Fix** using a low degree polynomial, so our simulation inherits the failure probability of **Fix** proved in [25] (and our simulation incurs only $\exp(-\Omega(n^{1/3}))$ additional error probability). We believe this failure probability can be improved to exponentially small, because the setting of [25] required the algorithm to find an *exact* satisfying assignment, whereas our notion of (δ, γ, η) -solve allows a small constant fraction of mistakes.

The requirement that $\eta > k^{-12}$ is likely also an artifact of our methods. This lower bound on the error tolerance arises because we simulate only the first phase of **Fix**. This phase produces an assignment within Hamming distance k^{-12} of a satisfying assignment, which is subsequently repaired by the rest of the algorithm. Because the solution geometry at clause density $(1 - \varepsilon)2^k \log k/k$ exhibits only short-range correlations, we believe it is possible to simulate **Fix** by a low degree polynomial (in fact, by a local algorithm, which we then simulate with a low degree polynomial as in Section 7.3) to arbitrary error tolerance, which would show Theorem 2.10 for any $\eta > 0$. We do not attempt these improvements in this paper.

Remark 2.12. By similar methods to our proof, the sequential local algorithms considered in [46] can also be simulated by a low degree polynomial. We will explain this in Section 7.2.

3 The Overlap Gap Program and Sketch of Main Ideas

In this section, we will discuss the recent line of work on the overlap gap property and place our methods in this context. The OGP program draws a rigorous connection between the clustering phenomenon and algorithmic hardness, and is the first line of work that translates properties of solution geometry to rigorous hardness results against classes of algorithms. OGP formalizes the clustering phenomenon as a prohibition of structures with medium overlaps, and then shows that any sufficiently smooth algorithm solving the problem can be used to construct the forbidden structure. Thus, such algorithms do not exist beyond the onset of (the appropriate notion of) clustering.

3.1 OGP for Maximum Independent Set

The first problem where a sharp algorithmic phase transition was derived using OGP was maximum independent set in sparse random graphs. In this problem, we are given a sample $G \sim G(n, d/n)$ of a sparse Erdős-Rényi graph and our task is to output an independent set of a specified size; the desired size of the set controls the problem difficulty. We work in the double limit where $n \rightarrow \infty$, and then $d \rightarrow \infty$. It is known [38, 10] that the largest independent set of this graph has size $\frac{2 \log d}{d} n$. More precisely, if $S_{\max}$ is a largest independent set, then as $n \rightarrow \infty$ for fixed d we have $\frac{1}{n} |S_{\max}| \rightarrow \alpha_d$, for some $\alpha_d = (1 + o_d(1)) \frac{2 \log d}{d}$. However, the best polynomial-time algorithm to date [57] can only find an independent set of asymptotic size $\frac{\log d}{d} n$, half the optimum. It is conjectured that no polynomial-time algorithm can find an independent set that is asymptotically larger. Rigorous results about this problem's solution geometry support this conjecture: Coja-Oghlan and Efthymiou [26] showed that independent sets of size $(1 + \varepsilon) \frac{\log d}{d} n$, for any constant $\varepsilon > 0$, are clustered in a way that implies slow mixing of any local Markov chain to sample these sets (but not necessarily that a local Markov chain cannot find a single such set).

Hardness against local algorithms via OGP and multi-OGP. The first negative result against a class of algorithms was proved by Gamarnik and Sudan [45], who showed that *local algorithms* (also called *factors of i.i.d.* algorithms) cannot find independent sets of asymptotic size larger than $\left(1 + \frac{1}{\sqrt{2}}\right) \frac{\log d}{d} n$. Their argument consists of two parts. First, they show that with high probability, a graph $G \sim G(n, d/n)$ does not have two independent sets of size asymptotically larger than $\left(1 + \frac{1}{\sqrt{2}}\right) \frac{\log d}{d} n$ with intersection in an interval $\left[(1 - \varepsilon) \frac{\log d}{d} n, (1 + \varepsilon) \frac{\log d}{d} n\right]$. Then, they construct an interpolation between two executions of a putative local algorithm that outputs an independent set of this size. The two executions are correlated, with some internal randomness shared and some internal randomness independent. By continuously tuning the amount of shared internal randomness, they extract from this interpolation two large independent sets with medium intersection, yielding a contradiction.

Rahman and Virág [69] generalized this argument, showing that for any $\varepsilon > 0$, local algorithms can find an independent set of size $(1 - \varepsilon) \frac{\log d}{d} n$, but not one of size $(1 + \varepsilon) \frac{\log d}{d} n$. The key idea of their negative result is to consider a forbidden overlap structure involving *several* large independent sets, generated from several correlated runs of a local algorithm. This is the first instance of a multi-OGP identifying a sharp algorithmic phase transition against some class of algorithms.

Hardness against low degree polynomials by the ensemble innovation. A later line of work extended this impossibility result to low degree polynomials, a much more powerful class of algorithms. Gamarnik, Jagannath, and Wein [43] showed that low degree polynomials cannot find independent sets of asymptotic size larger than $\left(1 + \frac{1}{\sqrt{2}}\right) \frac{\log d}{d} n$.² Their argument leverages an *ensemble OGP*, an idea introduced in [21]. In this approach, they construct an interpolation, this time over a sequence of *correlated problem instances*, whose endpoints are independent problem instances. They show that with high probability, there do not exist two independent sets, *possibly of different problem instances in the sequence*, of asymptotic size larger than $\left(1 + \frac{1}{\sqrt{2}}\right) \frac{\log d}{d} n$ with intersection in $\left[(1 - \varepsilon) \frac{\log d}{d} n, (1 + \varepsilon) \frac{\log d}{d} n\right]$. Due to the smoothness of low

²In this paper, they also show that the Hamiltonian of the spherical or Ising p -spin glass model cannot be optimized within some $\varepsilon > 0$ of its maximum by low degree polynomials or Langevin dynamics (applicable only to the spherical model).

degree polynomials, the outputs of a low degree polynomial on two consecutive instances of this interpolation are close with nontrivial probability. So, a putative low degree polynomial finding independent sets of this size can be used to construct such a forbidden structure with nontrivial probability, yielding a contradiction.

Wein [71] generalized this result by leveraging an *ensemble multi-OGP*, an approach that combines the multi-OGP and ensemble OGP ideas. In this approach, the interpolation is over a sequence of correlated problem instances, and the forbidden structure consists of several independent sets, possibly of different problem instances in the sequence, with prescribed multi-way overlaps. For any $\varepsilon > 0$, Wein showed that a putative low degree polynomial algorithm that finds independent sets of size at least $(1 + \varepsilon) \frac{\log d}{d} n$ can be used to construct the forbidden structure, and thus low degree polynomial algorithms cannot find independent sets of this size. Conversely, the local algorithms finding independent sets of size $(1 - \varepsilon) \frac{\log d}{d} n$ can be simulated by low degree polynomials. Thus, low degree polynomial algorithms find independent sets of asymptotic size $\frac{\log d}{d} n$ and no more.

Main idea of ensemble multi-OGP. At a high level, Wein's ensemble multi-OGP technique chains together many small negative free energy contributions to force a free energy to be negative. To simplify the discussion, we assume all the independent sets in the forbidden structure are independent sets of the same problem instance; this will capture the correct exponential rate (see Remark 5.4). Consider the normalized log first moment

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E}_{G \sim G(n, d/n)} \# \left(\begin{array}{c} (S^{(1)}, \dots, S^{(L)}) : S^{(1)}, \dots, S^{(L)} \text{ are independent} \\ \text{sets of } G \text{ of size } (1 + \varepsilon) \frac{\log d}{d} n \text{ satisfying } P \end{array} \right), \quad (3.1)$$

where P is a set of conditions on how $S^{(1)}, \dots, S^{(L)}$ overlap. This can be thought of as the free energy density of the uniform model over tuples $(S^{(1)}, \dots, S^{(L)})$ of independent sets of G of size $(1 + \varepsilon) \frac{\log d}{d} n$ satisfying P ; we will henceforth refer to (3.1) as a free energy. The structure inside the expectation in (3.1) is the forbidden structure we wish to rule out, defined in terms of overlaps between the sets. If the free energy (3.1) is negative, then with high probability this structure does not occur, and we say that the multi-OGP occurs.

The key idea is to set P as the intersection of conditions $P_2 \cap P_3 \cap \dots \cap P_L$, where P_ℓ is a condition on how $S^{(\ell)}$ overlaps with $S^{(1)}, \dots, S^{(\ell-1)}$, such that the following occurs for all $2 \leq \ell \leq L$.

- (1) Let $\mathcal{E}_\ell$ denote (3.1) with $(S^{(1)}, \dots, S^{(\ell)})$ in place of $(S^{(1)}, \dots, S^{(L)})$ and $P_2 \cap \dots \cap P_\ell$ in place of P . Then, $\mathcal{E}_\ell$ is smaller than $\mathcal{E}_{\ell-1}$ by an amount bounded away from 0. Informally, P_ℓ requires $S^{(\ell)}$ to overlap with its predecessors in a way that contributes a small negative free energy to (3.1).
- (2) For any fixed $S^{(1)}, \dots, S^{(\ell-1)}$, if $S^{(\ell)}$ starts at $S^{(\ell-1)}$, evolves without large jumps, and eventually evolves far away from all of $S^{(1)}, \dots, S^{(\ell-1)}$, then at some point along this evolution the condition P_ℓ occurs. Informally, the condition P_ℓ defines a high-dimensional moat that a stably evolving $S^{(\ell)}$ must cross.

Due to condition (1), if we set L large enough, (3.1) becomes negative, and the structure in (3.1) with high probability does not occur. Due to condition (2), we can construct this structure by taking L sets from a sequence of low degree polynomial outputs on correlated problem instances, which is (with nontrivial probability) a stable sequence. Specifically, we take $S^{(1)}$ to be the first output in the sequence, and then for $\ell \geq 2$ we take $S^{(\ell)}$ to be the first output after $S^{(\ell-1)}$ such that P_ℓ holds. This gets the desired contradiction.

The main technical challenge of this approach is to design the conditions P_ℓ such that (1) and (2) both hold. Effectively, one needs to construct a moat topologically disconnecting a high-dimensional space such that, for all values of $S^{(\ell)}$ in the moat, the free energy decrease required by (1) occurs. The requirement that the moat topologically disconnects the space gives us little control, and therein lies the difficulty.

In [71, Proposition 2.3], Wein carries out this approach by constructing P_ℓ as the condition that

$$\left| S^{(\ell)} \setminus (S^{(1)} \cup \dots \cup S^{(\ell-1)}) \right| \in \left[\frac{\varepsilon \log d}{4d} n, \frac{\varepsilon \log d}{2d} n \right]$$

and proving it has the required properties. Maximum independent set is amenable to the negative free energy chaining approach because the energy (3.1) enjoys a *geometric independence property*: given the induced subgraph of $G \sim G(n, d/n)$ on any subset of vertices $V \subseteq [n]$, the edges from V to any vertex

$v \in [n] \setminus V$ are independent randomness. This property makes the analysis of the free energy (3.1) tractable and shows in the relative simplicity of the definitions of the moats P_ℓ , which only consider each independent set's non-intersection with the union of its predecessors.

3.2 OGP for Random k -SAT and Our Contributions

This paper completes a similar picture for random k -SAT. Prior to our work, Gamarnik and Sudan [46] showed that balanced sequential local algorithms do not solve random NAE- k -SAT beyond clause density $(1 + o_k(1))2^k \log^2 k/k$ using a (non-ensemble) multi-OGP. They considered a forbidden overlap structure involving several satisfying assignments, generated from correlated runs of such an algorithm. They required the algorithm to be *balanced*, meaning that on any input, over the algorithm's internal randomness each of its outputs is unbiased; because of this requirement, their result required the symmetry provided by the NAE variant of random k -SAT. Balance is necessary in their interpolation argument to ensure that two fully independent outputs of the sequential local algorithm have high disagreement. We improve on this result in three ways:

- (1) We improve the threshold clause density by a logarithmic factor, to $(1 + o_k(1))\kappa^* 2^k \log k/k$.
- (2) We generalize the algorithm class ruled out from balanced sequential local algorithms to low degree polynomial algorithms.
- (3) We show hardness for random k -SAT instead of NAE- k -SAT; a simple adaptation of our argument shows hardness of random NAE- k -SAT at clause density $(1 + o_k(1))\kappa^* 2^{k-1} \log k/k$.

Improvements due to ensemble OGP. Unlike [46], we consider an ensemble multi-OGP, where the random variable being resampled is the k -SAT instance instead of the algorithm's internal randomness. This generalization allows us to rule out low degree polynomial algorithms, a much more powerful class than balanced sequential local algorithms. It also obviates the requirement of balance, so we no longer require the additional symmetry provided by NAE- k -SAT. This achieves improvements (2) and (3).

A tighter free energy analysis. Most crucially, we conduct a tighter free energy analysis than [46], which allows us to achieve improvement (1). In contrast to [46], whose forbidden structure only considers pairwise Hamming distances, the forbidden structure we use considers all 2^k ways $k+1$ satisfying assignments $y^{(0)}, \dots, y^{(k)}$ can agree or disagree. We formalize such an agreement pattern as an *overlap profile* π , which we rigorously introduce in Section 4.3. As shown in Lemma 5.1, the analogue of the free energy (3.1) for random k -SAT at clause density α for a set of overlap constraints P can then be expressed as

$$\log 2 + \max_{\pi \in P} \left[H(\pi) - \frac{\alpha}{2^k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right| \right],$$

where $y^{(0)}, \dots, y^{(k)}$ have overlap profile π and $\pi \in P$ denotes the set of overlap profiles π consistent with the constraints P . Here, $y^{(\ell)}[I]$ is the bit string obtained by indexing $y^{(\ell)}$ in positions I , namely $(y_{I_1}^{(\ell)}, \dots, y_{I_k}^{(\ell)})$. $H(\pi)$ is the entropy of a certain probability distribution associated to overlap profile π and the negative term is interpreted as an energy and captures the log likelihood that a random formula will have the $y^{(\ell)}$ as solutions. Whenever this free energy is negative, existence of tuples of satisfying assignments with profile $\pi \in P$ has exponentially small probability. We will choose P to capture an appropriate notion of (violation of) OGP.

Like before, we will chain together many small negative free energies to make this free energy negative. Unlike for maximum independent set, where there is the geometric independence property, for random k -SAT this free energy is highly dependent and hard to control. This makes it difficult to identify the correct high-dimensional moats in the negative free energy chaining approach. In the multi-OGP of [46], the condition P stipulates that the normalized Hamming distances $\Delta(y^{(i)}, y^{(j)})$ for the satisfying assignments in the forbidden structure are pairwise approximately $\frac{\log k}{k}$. Given this, the term $\mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right|$ in the free energy can be lower bounded by an inclusion-exclusion expansion truncated at level 2. This inclusion-exclusion estimate is not sharp, and consequently requires the larger value $\alpha = (1 + o_k(1))2^k \log^2 k/k$ to show

that the contribution of each $y^{(\ell)}$ to the free energy is a small negative number. The fact that this natural estimate of the free energy gives a threshold too large by a $\log k$ factor highlights the difficulty of accurately controlling the k -SAT free energy and the necessity of finding the correct high-dimensional moats.

We find the correct moats. Like above, we set $P = P_1 \cap \dots \cap P_k$, where P_ℓ governs how $y^{(\ell)}$ overlaps with its predecessors $y^{(0)}, \dots, y^{(\ell-1)}$ and defines a moat that a smooth evolution of $y^{(\ell)}$ starting from $y^{(\ell-1)}$ must cross. In order to obtain a fine control over the tradeoff between entropy and energy in the free energy, we develop a notion of *conditional overlap entropy* $H(\pi(y^{(\ell)}|y^{(0)}, \dots, y^{(\ell-1)}))$, which we think of as the contribution of $y^{(\ell)}$ to the entropy term $H(\pi)$. Our condition P_ℓ will stipulate that this conditional overlap entropy lies in an interval $\left[\beta_- \frac{\log k}{k}, \beta_+ \frac{\log k}{k}\right]$. This choice of forbidden structure *in terms of the conditional overlap entropy* is an important contribution of our work. The choice is motivated by the subsequent free energy analysis, which entails showing a lower bound on the energy contribution of each $y^{(\ell)}$ that counterbalances the entropy increase. We next summarize the argument.

Energy increment bound via decoupling. We can express the energy term as

$$\mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right| = \sum_{\sigma \in \{\mathbf{T}, \mathbf{F}\}^k} \mathbb{P}_{I \sim \text{unif}([n]^k)} \left[\sigma = y^{(\ell)}[I] \text{ for some } 0 \leq \ell \leq k \right]. \quad (3.2)$$

The complement of a probability on the right of (3.2) is the probability that $\sigma \neq y^{(\ell)}[I]$ for all $0 \leq \ell \leq k$. This can be conditionally expanded as a product of k factors, where the ℓ th factor is the probability that $\sigma \neq y^{(\ell)}[I]$ given the values of $y^{(0)}[I], \dots, y^{(\ell-1)}[I]$; we think of this factor as the contribution of $y^{(\ell)}$. We truncate these factors by rounding any that are less than $1 - \frac{1}{k \log k}$ up to 1. We choose the value $\frac{1}{k \log k}$ in order to apply the estimate that for $0 \leq \varepsilon_1, \dots, \varepsilon_k \leq \frac{1}{k \log k}$,

$$1 - (1 - \varepsilon_1)(1 - \varepsilon_2) \dots (1 - \varepsilon_k) \approx \varepsilon_1 + \varepsilon_2 + \dots + \varepsilon_k.$$

This decouples the contributions of the $y^{(\ell)}$. We can then bound the total contribution of $y^{(\ell)}$ to the left-hand side of (3.2) by summing the now-decoupled contributions over $\sigma \in \{\mathbf{T}, \mathbf{F}\}^k$.

Probabilistic reinterpretation. Miraculously, this sum can be reinterpreted as the success probability of an experiment involving a sum of k i.i.d. random variables, which can be controlled by concentration inequalities. We find that if the contribution of $y^{(\ell)}$ to $H(\pi)$ is $\beta \frac{\log k}{k}$, then its contribution to $\mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right|$ is at least $1 - \beta e^{-(\beta-1)}$. This motivates the choice of $\iota(\beta)$ as the ratio of these contributions, and κ^* as the best possible ratio. When $\alpha = \kappa 2^k \log k / k$ for $\kappa > \kappa^*$, the condition P_ℓ requires β to be in the range where the overall contribution of $y^{(\ell)}$ to the free energy is negative. Then, by the negative free energy chaining technique, the multi-OGP occurs.

This energy analysis via decoupling and probabilistic reinterpretation is original and is another key contribution of our work.

Future directions. Because we successfully implement the negative free energy chaining method for the random k -SAT free energy at an asymptotically optimal clause density, we believe our methods are evidence that the negative free energy chaining method can be used to establish sharp algorithmic hardness thresholds via multi-OGP for many other problems.

4 Forbidden Structures from Overlap Gap Property

This section and the next two sections are devoted to proving our main impossibility result, Theorem 2.6. Throughout these sections, we fix $\kappa > \kappa^*$. We set $\alpha = \kappa 2^k \log k / k$ and $m = \lfloor \alpha n \rfloor$.

4.1 Reduction to Deterministic Low Degree Polynomial

The following lemma shows that randomness does not significantly improve the power of low degree polynomial algorithms.

Lemma 4.1. *Suppose there exists a random degree- D polynomial that (δ, γ, η) -solves $\Phi_k(n, m)$. Then, there exists a deterministic degree- D polynomial that $(3\delta, 3\gamma, \eta)$ -solves $\Phi_k(n, m)$.*

Proof. Let f be a random degree- D polynomial that (δ, γ, η) -solves $\Phi_k(n, m)$. By definition of (δ, γ, η) -solve, we have $\mathbb{E}_\omega \mathbb{P}_\Phi [\mathcal{A}(\Phi, \omega) \text{ does not satisfy } \Phi] \leq \delta$ and $\mathbb{E}_\omega \mathbb{E}_\Phi [\|f(\Phi, \omega)\|_2^2] \leq \gamma n$. By Markov's inequality,

$$\mathbb{P}_\omega \left[\mathbb{P}_\Phi [\mathcal{A}(\Phi, \omega) \text{ does not satisfy } \Phi] \geq 3\delta \right] \leq \frac{1}{3} \quad \text{and} \quad \mathbb{P}_\omega \left[\mathbb{E}_\Phi [\|f(\Phi, \omega)\|_2^2] \geq 3\gamma n \right] \leq \frac{1}{3}.$$

So, there exists some ω such that $\mathbb{P}_\Phi [\mathcal{A}(\Phi, \omega) \text{ satisfies } \Phi] \geq 1 - 3\delta$ and $\mathbb{E}_\Phi [\|f(\Phi, \omega)\|_2^2] \leq 3\gamma n$. $\square$

By Lemma 4.1, it suffices to rule out deterministic low degree polynomials. For the rest of this section and Section 6, except where stated, $f : \mathbb{R}^N \rightarrow \mathbb{R}^n$ will be a deterministic degree- D polynomial and $\mathcal{A}$ is some η -assisted rounding of f , for η we will set later. Because f is deterministic, $\mathcal{A}$ is also deterministic.

4.2 The Interpolation Path

We can enumerate the km literals of a formula $\Phi \in \Omega_k(n, m)$ in lexicographic order:

$$\Phi_{1,1}, \Phi_{1,2}, \dots, \Phi_{1,k}, \Phi_{2,1}, \dots, \Phi_{m,k}.$$

For $1 \leq s \leq km$, let $L(\Phi, s)$ denote the s th literal in this order. Formally, $L(\Phi, s) = \Phi_{i,j}$, where (i, j) is the unique pair of integers satisfying $1 \leq i \leq m$, $1 \leq j \leq k$, and $k(i-1) + j = s$. We now define a sequence of correlated random k -SAT formulas that will be central to our argument.

Definition 4.2 (Interpolation path). Let $T = k^2m$. Let $\Phi^{(0)}, \dots, \Phi^{(T)} \in \Omega_k(n, m)$ be the interpolation path of k -SAT instances sampled as follows. First, sample $\Phi^{(0)} \sim \Phi_k(n, m)$. For each $1 \leq t \leq T$, let $\sigma(t)$ be the unique integer such that $1 \leq \sigma(t) \leq km$ and $t - \sigma(t)$ is a multiple of km . Then, $\Phi^{(t)}$ is obtained from $\Phi^{(t-1)}$ by resampling $L(\Phi^{(t)}, \sigma(t))$ from $\text{unif}(\mathcal{L})$. Moreover, for $0 \leq t \leq T$, let $x^{(t)} = \mathcal{A}(\Phi^{(t)})$.

In other words, we start from a random k -SAT instance and resample the literals one by one in lexicographic order. After we have resampled all the literals we start over, repeating the procedure until each literal has been resampled k times. Note that each $\Phi^{(t)}$ is marginally a sample from $\Phi_k(n, m)$. Further, note that if $|t - t'| \geq km$, then $\Phi^{(t)}$ and $\Phi^{(t')}$ are independent. We run our assisted low degree algorithm $\mathcal{A}$ on all these k -SAT instances and collect the outputs as the sequence $x^{(0)}, x^{(1)}, \dots, x^{(T)}$.

4.3 Overlap Profiles

We introduce the notion of the overlap profile of an ordered list of assignments, which will be central to our proof of impossibility. The overlap profile summarizes the bitwise agreement and disagreement pattern of a list of assignments. In the proof of Theorem 2.6, we will use these overlap profiles to reason about the assignments $x^{(0)}, \dots, x^{(T)}$ arising in the interpolation of Definition 4.2.

We begin with a notion of partition. Let $\mathcal{P}_2(\ell)$ denote the set of unordered partitions of $\{0, \dots, \ell-1\}$ into two (possibly empty) sets. For example, the set $\mathcal{P}_2(3)$ consists of the following four partitions:

$$\{\{0, 1, 2\}, \emptyset\}, \{\{0, 1\}, \{2\}\}, \{\{0, 2\}, \{1\}\}, \{\{1, 2\}, \{0\}\}.$$

Note that $|\mathcal{P}_2(\ell)| = 2^{\ell-1}$.

Definition 4.3 (Overlap profile). Let $y^{(0)}, \dots, y^{(\ell-1)} \in \{\mathbf{T}, \mathbf{F}\}^n$ be a list of ℓ assignments. The *overlap profile* of $y^{(0)}, \dots, y^{(\ell-1)}$, denoted $\pi(y^{(0)}, \dots, y^{(\ell-1)})$, is a vector $\pi \in \mathbb{R}^{2^{\ell-1}}$ indexed by unordered pairs $\{S, T\} \in \mathcal{P}_2(\ell)$, where

$$\pi_{S,T} = \frac{1}{n} \left| \left\{ i \in [n] : \text{all } \{y_i^{(t)} : t \in S\} \text{ equal one value and all } \{y_i^{(t)} : t \in T\} \text{ equal the other value} \right\} \right|.$$

Example 4.4. Let $\ell = 3$. The overlap profile $\pi = \pi(y^{(0)}, y^{(1)}, y^{(2)})$ consists of four entries $\pi_{012, \emptyset}$, $\pi_{01, 2}$, $\pi_{02, 1}$, and $\pi_{12, 0}$, where

$$\pi_{012, \emptyset} = \frac{1}{n} \left| i \in [n] : y_i^{(0)} = y_i^{(1)} = y_i^{(2)} \right| \quad \text{and} \quad \pi_{01, 2} = \frac{1}{n} \left| i \in [n] : y_i^{(0)} = y_i^{(1)} \neq y_i^{(2)} \right|,$$

and analogously for $\pi_{02, 1}$ and $\pi_{12, 0}$.

We can understand an overlap profile as a probability distribution: $\pi_{S, T}$ is the probability that in a uniformly random position $i \sim \text{unif}([n])$, all the $\{y_i^{(t)} : t \in S\}$ agree, all the $\{y_i^{(t)} : t \in T\}$ agree, and these two sets disagree with each other. Thus, we can naturally define the entropy of an overlap profile by

$$H\left(\pi(y^{(0)}, \dots, y^{(\ell-1)})\right) = - \sum_{\{S, T\} \in \mathcal{P}_2(\ell)} \pi_{S, T} \log \pi_{S, T}.$$

This is the entropy of the unordered pair of sets $\{S, T\}$ obtained by sampling $i \sim \text{unif}([n])$ and partitioning $\{0, \dots, \ell-1\}$ based on the value of $y_i^{(t)}$.

As the overlap profile of $y^{(0)}, \dots, y^{(t-1)}$ refines the overlap profile of any subsequence of $y^{(0)}, \dots, y^{(t-1)}$, we can extend Definition 4.3 verbatim to unordered pairs $\{S, T\}$ of disjoint sets $S, T \subseteq \{0, \dots, \ell-1\}$. These overlaps are determined by the overlaps $\pi_{S, T}$ for $\{S, T\} \in \mathcal{P}_2(\ell)$: for example, when $\ell = 3$, $\pi_{0, 1} = \pi_{02, 1} + \pi_{12, 0}$.

This refinement property also allows us to define a notion of conditional overlap profile. Let $\pi = \pi(y^{(0)}, \dots, y^{(\ell-1)})$. Then, we can define the conditional overlap profile $\pi_{\cdot|S} = \pi(y^{(\ell-1)} | y^{(0)}, \dots, y^{(\ell-2)})$ as follows. For each $\{S, T\} \in \mathcal{P}_2(\ell-1)$ with $\pi_{S, T} > 0$, $\pi_{\cdot|S, T}$ is a probability distribution on the two partitions $\{S \cup \{\ell-1\}, T\}$ and $\{S, T \cup \{\ell-1\}\}$ with

$$\pi_{S \cup \{\ell-1\}, T | S, T} = \frac{\pi_{S \cup \{\ell-1\}, T}}{\pi_{S, T}} \quad \text{and} \quad \pi_{S, T \cup \{\ell-1\} | S, T} = \frac{\pi_{S, T \cup \{\ell-1\}}}{\pi_{S, T}}.$$

(If $\pi_{S, T} = 0$, we can define this distribution arbitrarily.) This is the probability distribution of the agreement pattern of $y^{(0)}, \dots, y^{(\ell-1)}$ on a uniformly random position, conditioned on the agreement pattern of $y^{(0)}, \dots, y^{(\ell-2)}$ in that position being $\{S, T\}$. We can analogously define the conditional overlap entropy by

$$H\left(\pi(y^{(\ell-1)} | y^{(0)}, \dots, y^{(\ell-2)})\right) = \sum_{\{S, T\} \in \mathcal{P}_2(\ell-1)} \pi_{S, T} H(\pi_{\cdot|S, T}).$$

Before proceeding, we collect some properties of overlap profiles which will be useful in the rest of the section. The proofs of these assertions follow readily from the above definitions.

Fact 4.5. *Overlap profiles have the following properties.*

- (a) *There are at most $n^{2^{\ell-1}}$ distinct overlap profiles of ℓ assignments $y^{(0)}, \dots, y^{(\ell-1)} \in \{\mathbf{T}, \mathbf{F}\}^n$.*
- (b) *The entropy of overlap profiles satisfies the chain rule*

$$H\left(\pi(y^{(0)}, \dots, y^{(\ell-1)})\right) = H\left(\pi(y^{(0)}, \dots, y^{(\ell-2)})\right) + H\left(\pi(y^{(\ell-1)} | y^{(0)}, \dots, y^{(\ell-2)})\right).$$

- (c) *If $z^{(0)}, \dots, z^{(r-1)}$ are the distinct elements of $y^{(0)}, \dots, y^{(\ell-1)}$, then*

$$H\left(\pi(y^{(0)}, \dots, y^{(\ell-1)})\right) = H\left(\pi(z^{(0)}, \dots, z^{(r-1)})\right).$$

If $z^{(0)}, \dots, z^{(r-2)}$ are the distinct elements of $y^{(0)}, \dots, y^{(\ell-2)}$, then

$$H\left(\pi(y^{(\ell-1)} | y^{(0)}, \dots, y^{(\ell-2)})\right) = H\left(\pi(y^{(\ell-1)} | z^{(0)}, \dots, z^{(r-2)})\right).$$

Furthermore, if $y^{(\ell-1)} \in \{y^{(0)}, \dots, y^{(\ell-2)}\}$, then $H\left(\pi(y^{(\ell-1)} | y^{(0)}, \dots, y^{(\ell-2)})\right) = 0$.

4.4 Outline of Proof of Impossibility

Recall that $\iota(\beta) = \frac{\beta}{1-\beta e^{-(\beta-1)}}$ is convex with minimum κ^* attained at β^* . This function is strictly decreasing on $(1, \beta^*]$ and strictly increasing on $[\beta^*, +\infty)$. Moreover, $\iota(\beta)$ tends to $+\infty$ when $\beta \rightarrow 1^+$ or $\beta \rightarrow +\infty$.

Because $\kappa > \kappa^*$, there exist two solutions $\beta_{\min}, \beta_{\max}$ to $\iota(\beta) = \kappa$, with $\beta_{\min} \in (1, \beta^*)$ and $\beta_{\max} \in (\beta^*, +\infty)$. Set $\beta_- = \frac{\beta_{\min} + \beta^*}{2}$ and $\beta_+ = \frac{\beta_{\max} + \beta^*}{2}$. (This choice is arbitrary; any deterministic choice with $\beta_{\min} < \beta_- < \beta_+ < \beta_{\max}$ will do.) Set $\varepsilon > 0$ such that $\frac{\beta_+ + \varepsilon}{1 - \beta_+ e^{-(\beta_+ - 1)}} \leq \kappa$ for all $\beta \in [\beta_-, \beta_+]$. Thus $\beta_-, \beta_+, \varepsilon$ are all deterministic functions of κ .

We will define the events $S_{\text{valid}}, S_{\text{consec}}, S_{\text{indep}}, S_{\text{ogp}}$ as the following events of the interpolation path defined in Definition 4.2. Define

$$S_{\text{valid}} = \left\{ x^{(t)} \text{ satisfies } \Phi^{(t)} \text{ for all } 0 \leq t \leq T \right\}.$$

This is the event that $\mathcal{A}$ succeeds on all instances in the interpolation path. Define

$$S_{\text{consec}} = \left\{ \Delta(x^{(t)}, x^{(t-1)}) \leq \frac{\beta_+ - \beta_-}{2k} \text{ for all } 1 \leq t \leq T \right\}.$$

This is the event that outputs of $\mathcal{A}$ on consecutive instances in the interpolation path are not too far in Hamming distance. Define S_{indep} as the event that there do not exist $k+1$ indices $0 \leq t_0 \leq t_1 \leq \dots \leq t_k \leq T$ and an assignment $y \in \{\mathbf{T}, \mathbf{F}\}^n$ such that

(IND-A) For all $0 \leq \ell \leq k-1$, $|t_k - t_\ell| \geq km$;

(IND-B) y satisfies $\Phi^{(t_k)}$; and

(IND-C) $H(\pi(y|x^{(t_0)}, \dots, x^{(t_{k-1})})) < \beta_- \frac{\log k}{k}$.

This is the event that relative to any collection of outputs of $\mathcal{A}$, all the solutions to an *independent* k -SAT instance have high conditional overlap entropy. Finally, define S_{ogp} as the event that there do not exist $k+1$ indices $0 \leq t_0 \leq t_1 \leq \dots \leq t_k \leq T$ and assignments $y^{(0)}, \dots, y^{(k)} \in \{\mathbf{T}, \mathbf{F}\}^n$ such that

(OGP-A) For all $0 \leq \ell \leq k$, $y^{(\ell)}$ satisfies $\Phi^{(t_\ell)}$; and

(OGP-B) For all $1 \leq \ell \leq k$, $H(\pi(y^{(\ell)}|y^{(0)}, \dots, y^{(\ell-1)})) \in \left[\beta_- \frac{\log k}{k}, \beta_+ \frac{\log k}{k} \right]$.

The structure ruled out by S_{ogp} is the main forbidden structure of our argument. Informally, this forbidden structure consists of $k+1$ assignments, satisfying possibly different k -SAT instances in the interpolation path, where each satisfying assignment has a medium conditional overlap entropy relative to its predecessors.

The key ingredients in our proof of Theorem 2.6 are the following two propositions. Proposition 4.6 establishes the key relationship between the four events defined above. Proposition 4.7 controls the probabilities of these events. We will use these two propositions to derive the main contradiction in the proof of Theorem 2.6: if a low degree polynomial algorithm (δ, γ, η) solves $\Phi_k(n, m)$ for the requisite (δ, γ, η) , then Proposition 4.6 implies $S_{\text{valid}} \cap S_{\text{consec}} \cap S_{\text{indep}} \cap S_{\text{ogp}} = \emptyset$, while Proposition 4.7 implies $S_{\text{valid}} \cap S_{\text{consec}} \cap S_{\text{indep}} \cap S_{\text{ogp}} \neq \emptyset$.

Proposition 4.6. *For all sufficiently large k , $S_{\text{valid}} \cap S_{\text{consec}} \cap S_{\text{indep}} \subseteq S_{\text{ogp}}^c$.*

This proposition states that if S_{valid} , S_{consec} , and S_{indep} all occur, we can construct an example of the structure forbidden by S_{ogp} .

Proposition 4.7. *Suppose f is a deterministic degree- D polynomial that (δ, γ, η) -solves $\Phi_k(n, m)$, where $\eta = \frac{\beta_+ - \beta_-}{8k}$. For all sufficiently large k , the following inequalities hold.*

$$(a) \mathbb{P}(S_{\text{valid}} \cap S_{\text{consec}}) \geq (2n)^{-4\gamma D k^2 / (\beta_+ - \beta_-)} - (T+1)\delta.$$

$$(b) \mathbb{P}(S_{\text{indep}}^c) \leq \exp(-\Omega(n)).$$

$$(c) \mathbb{P}(S_{\text{ogp}}^c) \leq \exp(-\Omega(n)).$$

The remainder of this section and Sections 5 and 6 will be devoted to proving these two propositions. We will prove Proposition 4.6 in Section 4.5 and Proposition 4.7(b) in Section 4.6. We will prove Proposition 4.7(c), which establishes the presence of the main multi-OGP, in Section 5. Finally, we will prove Proposition 4.7(a) in Section 6. Let us first see how these results imply Theorem 2.6.

Proof of Theorem 2.6. Set $\eta = \frac{\beta_+ - \beta_-}{8k}$. Assume for sake of contradiction that there exists a (random) degree- D polynomial $g : \mathbb{R}^N \rightarrow \mathbb{R}^n$ that (δ, γ, η) -solves $\Phi_k(n, m)$. By Lemma 4.1, there exists a deterministic degree- D polynomial $f : \mathbb{R}^N \rightarrow \mathbb{R}^n$ that $(3\delta, 3\gamma, \eta)$ -solves $\Phi_k(n, m)$. We set $k^* = k^*(\kappa)$ large enough that Propositions 4.6 and 4.7 both hold.

By Proposition 4.6, $S_{\text{valid}} \cap S_{\text{consec}} \cap S_{\text{indep}} \cap S_{\text{ogp}} = \emptyset$. By Proposition 4.7,

$$\mathbb{P}(S_{\text{valid}} \cap S_{\text{consec}} \cap S_{\text{indep}} \cap S_{\text{ogp}}) \geq (2n)^{-12\gamma Dk^2/(\beta_+ - \beta_-)} - 3(T+1)\delta - \exp(-\Omega(n)).$$

We will show this probability is positive, which implies that $S_{\text{valid}} \cap S_{\text{consec}} \cap S_{\text{indep}} \cap S_{\text{ogp}} \neq \emptyset$ and yields a contradiction.

Pick $C_2 = 2 + \frac{12k^2}{\beta_+ - \beta_-}$. Recall that $T = k^2m = k^2\lfloor \alpha n \rfloor$. We can check that if $\delta \leq \exp(-C_2\gamma D \log n)$, then $3(T+1)\delta \leq \frac{1}{3}(2n)^{-12\gamma Dk^2/(\beta_+ - \beta_-)}$ for sufficiently large n . We can pick C_1 small enough that if $D \leq \frac{C_1 n}{\gamma \log n}$, then

$$(2n)^{-12\gamma Dk^2/(\beta_+ - \beta_-)} \geq n^{-24\gamma Dk^2/(\beta_+ - \beta_-)} \geq \exp\left(-\frac{24C_1k^2}{\beta_+ - \beta_-}n\right)$$

is asymptotically larger than the $\exp(-\Omega(n))$ term. Then, for sufficiently large n , the $\exp(-\Omega(n))$ term is at most $\frac{1}{3}(2n)^{-12\gamma Dk^2/(\beta_+ - \beta_-)}$. Therefore, there exists n^* such that if $n \geq n^*$, then $\mathbb{P}(S_{\text{valid}} \cap S_{\text{consec}} \cap S_{\text{indep}} \cap S_{\text{ogp}}) > 0$. $\square$

4.5 Constructing the Forbidden Structure from Low Degree Polynomial Outputs

In this section, we will prove Proposition 4.6, that if S_{valid} , S_{consec} , and S_{indep} all hold, then an instance of the structure forbidden by S_{ogp} exists.

We will need the following auxiliary lemma, which shows that a small change of $x \in \{\mathbf{T}, \mathbf{F}\}^n$ in Hamming distance induces a small change in the conditional overlap entropy $H(\pi(x|y^{(0)}, \dots, y^{(\ell-1)}))$. This lemma allows us to convert S_{consec} to a guarantee that consecutive conditional overlap entropies are small. We defer the proof of this lemma to Appendix A.

Lemma 4.8. *Let $\ell \in \mathbb{N}$ be arbitrary and let $x, x', y^{(0)}, \dots, y^{(\ell-1)} \in \{\mathbf{T}, \mathbf{F}\}^n$. If $\Delta(x, x') \leq \frac{1}{2}$, then*

$$\left| H\left(\pi(x|y^{(0)}, \dots, y^{(\ell-1)})\right) - H\left(\pi(x'|y^{(0)}, \dots, y^{(\ell-1)})\right) \right| \leq H(\Delta(x, x')).$$

The $H(\cdot)$ on the right denotes the binary entropy function.

Proof of Proposition 4.6. Set k large enough that $\frac{\beta_+ - \beta_-}{2k} \leq \frac{1}{2}$ and $H\left(\frac{\beta_+ - \beta_-}{2k}\right) \leq (\beta_+ - \beta_-)\frac{\log k}{k}$. The second inequality holds for all sufficiently large k due to the inequality $H(x) \leq x \log \frac{e}{x}$.

Suppose that S_{valid} , S_{consec} , and S_{indep} all hold. We will construct an example of the structure forbidden by S_{ogp} . For $0 \leq \ell \leq k$, we will set $y^{(\ell)} = x^{(t_\ell)}$, where $0 \leq t_0 \leq t_1 \leq \dots \leq t_k \leq T$ are defined as follows. Let $t_0 = 0$. For $1 \leq \ell \leq k$, let t_ℓ be the smallest $t > t_{\ell-1}$ such that $H(x^{(t)}|y^{(0)}, \dots, y^{(\ell-1)}) \in \left[\beta_- \frac{\log k}{k}, \beta_+ \frac{\log k}{k}\right]$. We will show that such t_ℓ exists and satisfies $t_\ell \leq t_{\ell-1} + km$.

At $t = t_{\ell-1}$, we have $H(\pi(x^{(t)}|y^{(0)}, \dots, y^{(\ell-1)})) = 0$ by Fact 4.5(c). Consider repeatedly incrementing t . Because S_{consec} holds, we have $\Delta(x^{(t)}, x^{(t-1)}) \leq \frac{\beta_+ - \beta_-}{2k}$. By Lemma 4.8, this implies that

$$\begin{aligned} \left| H\left(\pi(x^{(t)}|y^{(0)}, \dots, y^{(\ell-1)})\right) - H\left(\pi(x^{(t-1)}|y^{(0)}, \dots, y^{(\ell-1)})\right) \right| &\leq H(\Delta(x^{(t)}, x^{(t-1)})) \\ &\leq H\left(\frac{\beta_+ - \beta_-}{2k}\right) \leq (\beta_+ - \beta_-)\frac{\log k}{k}. \end{aligned}$$

Here, we use that $\frac{\beta_+ - \beta_-}{2k} \leq \frac{1}{2}$ and $H\left(\frac{\beta_+ - \beta_-}{2k}\right) \leq (\beta_+ - \beta_-) \frac{\log k}{k}$. Thus, $H(\pi(x^{(t)}|y^{(0)}, \dots, y^{(\ell-1)}))$ never skips over the interval $\left[\beta_- \frac{\log k}{k}, \beta_+ \frac{\log k}{k}\right]$. We will show that for $t' = t_{\ell-1} + km$, $H(\pi(x^{(t')}|y^{(0)}, \dots, y^{(\ell-1)})) \geq \beta_- \frac{\log k}{k}$; this will imply that $H(\pi(x^{(t)}|y^{(0)}, \dots, y^{(\ell-1)})) \in \left[\beta_- \frac{\log k}{k}, \beta_+ \frac{\log k}{k}\right]$ for some $t_{\ell-1} < t \leq t'$.

In the definition of S_{indep} , set $t_{\ell-1} = t_\ell = t_{\ell+1} = \dots = t_{k-1}$ and $t_k = t'$. Then (using Fact 4.5(c), which allows us to ignore the duplicated $t_\ell, \dots, t_{k-1}$), we see that $\Phi^{(t')}$ has no satisfying assignment y with $H(\pi(y|y^{(0)}, \dots, y^{(\ell-1)})) < \beta_- \frac{\log k}{k}$. But, because S_{valid} holds, $x^{(t')}$ satisfies $\Phi^{(t')}$. It follows that $H(\pi(x^{(t')}|y^{(0)}, \dots, y^{(\ell-1)})) \geq \beta_- \frac{\log k}{k}$.

Because the interpolation path has length $T = k^2 m$, and $t_\ell \leq t_{\ell-1} + km$ for all $1 \leq \ell \leq k$, this procedure sets all of $t_1, \dots, t_k$ before the end of the interpolation.

Finally, because S_{valid} holds, y_ℓ satisfies $\Phi^{(t_\ell)}$ for all $0 \leq \ell \leq k$. We have thus constructed the structure forbidden by S_{ogp} . $\square$

4.6 Solutions to Independent Instances Contribute Large Overlap Entropy

In this section, we will prove Proposition 4.7(b), that if $\Phi^{(t_k)}$ is independent of $x^{(t_0)}, \dots, x^{(t_{k-1})}$, then satisfying assignments y of $\Phi^{(t_k)}$ do not have conditional overlap entropy with $x^{(t_0)}, \dots, x^{(t_{k-1})}$ that is too small. We will prove this proposition by a first moment argument.

Proof of Proposition 4.7(b). By Markov's inequality, $\mathbb{P}(S_{\text{indep}}^c)$ is upper bounded by the expected number of $(t_0, \dots, t_k, y)$ satisfying conditions (IND-A), (IND-B), and (IND-C) of S_{indep} . There are at most $(T+1)^{k+1}$ possible choices of $0 \leq t_0 \leq t_1 \leq \dots \leq t_k \leq T$ satisfying condition (IND-A). By condition (IND-A), $\Phi^{(t_k)}$ is independent of $x^{(t_0)}, \dots, x^{(t_{k-1})}$.

Let $P = P(x^{(t_0)}, \dots, x^{(t_{k-1})})$ denote the set of all overlap profiles $\pi = \pi(x^{(t_0)}, \dots, x^{(t_{k-1})}, y)$ over $y \in \{\mathbf{T}, \mathbf{F}\}^n$ with $H(\pi(y|x^{(t_0)}, \dots, x^{(t_{k-1})})) < \beta_- \frac{\log k}{k}$. By Fact 4.5(a), $|P| \leq n^{2^k}$. Thus,

$$\begin{aligned} \mathbb{P}(S_{\text{indep}}^c) &\leq (T+1)^{k+1} n^{2^k} \max_{\substack{0 \leq t_0 \leq \dots \leq t_k \leq T \\ \text{satisfying (IND-A)}}} \max_{\pi \in P(x^{(t_0)}, \dots, x^{(t_{k-1})})} \\ &\quad \mathbb{E}_{\Phi^{(t_k)}} \left[\# \left(y \in \{\mathbf{T}, \mathbf{F}\}^n : y \text{ satisfies } \Phi^{(t_k)} \text{ and } \pi(x^{(t_0)}, \dots, x^{(t_{k-1})}, y) = \pi \right) \right] \end{aligned}$$

We can evaluate this inner expectation by linearity of expectation. The number of y satisfying that $\pi(x^{(t_0)}, \dots, x^{(t_{k-1})}, y) = \pi$ is

$$\begin{aligned} \prod_{\{S, T\} \in \mathcal{P}_2(k)} \binom{\pi_{S, T} n}{\pi_{S \cup \{k\}, T} n} &= \exp \left(n \sum_{\{S, T\} \in \mathcal{P}_2(k)} \pi_{S, T} H \left(\frac{\pi_{S \cup \{k\}, T}}{\pi_{S, T}} \right) + o(n) \right) \\ &= \exp \left(n H \left(\pi(y|x^{(t_0)}, \dots, x^{(t_{k-1})}) \right) + o(n) \right) \\ &\leq \exp \left(n \beta_- \frac{\log k}{k} + o(n) \right). \end{aligned}$$

Because $\Phi^{(t_k)}$ is independent of $x^{(t_0)}, \dots, x^{(t_{k-1})}$, the probability that any one of these y satisfies $\Phi^{(t_k)}$ is

$$(1 - 2^{-k})^m \leq \exp(-2^{-k} m) = \exp \left(-n \kappa \frac{\log k}{k} + o(n) \right).$$

Here we used that $m = \lfloor \alpha n \rfloor$ and $\alpha = \kappa 2^k \log k / k$. Thus,

$$\mathbb{P}(S_{\text{indep}}^c) \leq \exp \left(-n(\kappa - \beta_-) \frac{\log k}{k} + o(n) \right),$$

where the $(T+1)^{k+1} n^{2^k}$ is absorbed in the $o(n)$. Finally, note that $\beta_- + \varepsilon \leq \frac{\beta_- + \varepsilon}{1 - \beta_- e^{-(\beta_- - 1)}} \leq \kappa$, so $\kappa - \beta_- \geq \varepsilon$. Thus $\mathbb{P}(S_{\text{indep}}^c) = \exp(-\Omega(n))$. $\square$

5 Proof of Presence of Ensemble Multi-OGP

In this section, we will prove Proposition 4.7(c), which shows that the forbidden structure in S_{ogp} does not occur with high probability.

5.1 Proof Outline

We first give a high level overview of the proof of Proposition 4.7(c).

The proof is by another first moment computation. Throughout this section, for a k -tuple of indices $I \in [n]^k$ and $x \in \{\mathbf{T}, \mathbf{F}\}^n$, let $x[I] = (x_{I_1}, \dots, x_{I_k})$ be the string of bits in x indexed by I . We begin with the following lemma, which bounds the exponential rate of $\mathbb{P}(S_{\text{ogp}}^c)$ in terms of a maximum over overlap profiles. We will prove this lemma in Section 5.2.

Lemma 5.1. *Let P denote the set of overlap profiles $\pi = \pi(y^{(0)}, \dots, y^{(k)})$ over $y^{(0)}, \dots, y^{(k)} \in \{\mathbf{T}, \mathbf{F}\}^n$ satisfying that for all $1 \leq \ell \leq k$, $H(\pi(y^{(\ell)}|y^{(0)}, \dots, y^{(\ell-1)})) \in [\beta_- \frac{\log k}{k}, \beta_+ \frac{\log k}{k}]$. The following inequality holds.*

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(S_{\text{ogp}}^c) \leq \log 2 + \max_{\pi \in P} \left[H(\pi) - \kappa \frac{\log k}{k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right| \right], \quad (5.1)$$

where $y^{(0)}, \dots, y^{(k)} \in \{\mathbf{T}, \mathbf{F}\}^n$ is a sequence of assignments with overlap profile π .

Note that the expectation $\mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right|$ has the same value for any sequence of assignments $y^{(0)}, \dots, y^{(k)}$ with overlap profile π . So, the quantity inside the maximum is a function of π .

The negative term in (5.1) arises as an upper bound on the exponential rate of the probability that $y^{(0)}, \dots, y^{(k)}$ all respectively satisfy $\Phi^{(t_0)}, \dots, \Phi^{(t_k)}$, for fixed $y^{(0)}, \dots, y^{(k)}$ and $t_0, \dots, t_k$. Let us first argue heuristically that this bounds the exponential rate; we will formalize this reasoning in Lemma 5.3 below. We expect the probability that $y^{(0)}, \dots, y^{(k)}$ satisfy $\Phi^{(t_0)}, \dots, \Phi^{(t_k)}$ to be maximized when $t_0 = \dots = t_k$, because making the t_i different only introduces additional randomness (see Remark 5.4). So, let $\Phi^{(t_0)}, \dots, \Phi^{(t_k)}$ all equal the same k -SAT instance $\Phi \sim \Phi_k(n, m)$. The probability that $y^{(0)}, \dots, y^{(k)}$ all satisfy the first clause Φ_1 of Φ is $1 - 2^{-k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right|$, because if Φ_1 contains the variables $x_{I_1}, \dots, x_{I_k}$, there are exactly $\left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right|$ ways to set these variables' polarities in Φ_1 so that one of $y^{(0)}, \dots, y^{(k)}$ does not satisfy Φ_1 . Then, the probability that $y^{(0)}, \dots, y^{(k)}$ all satisfy Φ is upper bounded by

$$\left(1 - 2^{-k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right| \right)^m \approx \exp \left(-n \kappa \frac{\log k}{k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right| \right).$$

The second ingredient in the proof of Proposition 4.7(c) is the following proposition, which lower bounds the expectation in the negative term of (5.1). We will prove this proposition in Section 5.3. Proving the bound in this proposition is one of the main technical challenges of this paper, which we overcome via a surprising probabilistic reformulation of the left-hand expectation.

Proposition 5.2. *Let $\beta_1, \dots, \beta_k \in [\beta_-, \beta_+]$, and let $y^{(0)}, \dots, y^{(k)} \in \{\mathbf{T}, \mathbf{F}\}^n$ be assignments satisfying that $H(\pi(y^{(\ell)}|y^{(0)}, \dots, y^{(\ell-1)})) = \beta_\ell \frac{\log k}{k}$. Then,*

$$\mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right| \geq (1 - o_k(1)) \sum_{\ell=1}^k \left(1 - \beta_\ell e^{-(\beta_\ell - 1)} \right).$$

From Lemma 5.1 and Proposition 5.2, we can see the main ideas of the proof of Proposition 4.7(c) and understand the motivation of the definition of S_{ogp} . The main ideas are as follows.

We will prove Proposition 4.7(c) by showing that the right-hand side of (5.1) is negative. For each $\pi \in P$, this quantity can be regarded as a free energy, with entropy term $\log 2 + H(\pi)$ and energy term $\kappa \frac{\log k}{k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right|$. This free energy exhibits a tradeoff where as the entropy term increases, the assignments $y^{(0)}, \dots, y^{(k)}$ become more diverse, and so the energy term increases too. The event S_{ogp} is selected so that for overlap profiles $\pi \in P$, where P is defined in Lemma 5.1, the energy term is

larger than the entropy term, which makes the free energy negative. In particular, (due to Fact 4.5(b)) we think of $H(\pi(y^{(\ell)}|y^{(0)}, \dots, y^{(\ell-1)}))$ as the amount that $y^{(\ell)}$ contributes to the free energy's entropy term. Given this contribution, Proposition 5.2 lower bounds the amount that $y^{(\ell)}$ contributes to the energy term. In the definition of S_{ogp} , we require the entropy contribution to be in a medium range $[\beta_- \frac{\log k}{k}, \beta_+ \frac{\log k}{k}]$, because in this range the energy-to-entropy ratio is favorable to the energy term. Specifically, we show that if $y^{(\ell)}$ contributes an entropy in this range, the energy it contributes is at least $\varepsilon \frac{\log k}{k}$ more. Thus each $y^{(\ell)}$ decreases the free energy by at least $\varepsilon \frac{\log k}{k}$. Together, the k assignments $y^{(1)}, \dots, y^{(k)}$ contribute a free energy decrease of $\varepsilon \log k$, which dominates the starting free energy of $\log 2$ and makes the overall free energy negative.

We now formally prove Proposition 4.7(c) given Lemma 5.1 and Proposition 5.2.

Proof of Proposition 4.7(c). We begin from the bound (5.1). Let P be as in Lemma 5.1. Let $\pi \in P$, and let $y^{(0)}, \dots, y^{(k)} \in \{\mathbf{T}, \mathbf{F}\}^n$ with $\pi(y^{(0)}, \dots, y^{(k)}) = \pi$. Define $\beta_1, \dots, \beta_k$ as in Proposition 5.2; note that these are determined given π . By definition of P , we have $\beta_1, \dots, \beta_k \in [\beta_-, \beta_+]$. By Fact 4.5(b), $H(\pi) = \frac{\log k}{k} \sum_{\ell=1}^k \beta_\ell$. By Proposition 5.2,

$$\begin{aligned} -\kappa \frac{\log k}{k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right| &\leq -(1 - o_k(1)) \frac{\log k}{k} \sum_{\ell=1}^k \kappa \left(1 - \beta_\ell e^{-(\beta_\ell - 1)} \right) \\ &\leq -(1 - o_k(1)) \frac{\log k}{k} \sum_{\ell=1}^k (\beta_\ell + \varepsilon). \end{aligned}$$

The last inequality uses that $\frac{\beta + \varepsilon}{1 - \beta e^{-(\beta - 1)}} \leq \kappa$ for all $\beta \in [\beta_-, \beta_+]$. Therefore,

$$\begin{aligned} H(\pi) - \kappa \frac{\log k}{k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right| &\leq o_k(1) \frac{\log k}{k} \sum_{\ell=1}^k \beta_\ell - (1 - o_k(1)) \varepsilon \log k \\ &\leq o_k(1) \beta_+ \log k - (1 - o_k(1)) \varepsilon \log k \\ &= -(1 - o_k(1)) \varepsilon \log k. \end{aligned}$$

This bound holds for an arbitrary $\pi \in P$, and thus for the maximum over $\pi \in P$. So, by (5.1),

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \mathbb{P}(S_{\text{ogp}}^c) \leq \log 2 - (1 - o_k(1)) \varepsilon \log k < 0$$

for sufficiently large k . Thus $\mathbb{P}(S_{\text{ogp}}^c) \leq \exp(-\Omega(n))$. $\square$

5.2 Bounding the Exponential Rate by a Free Energy

In this section, we will prove Lemma 5.1. We begin with the following lemma, which bounds the probability term arising in the first moment upper bound of $\mathbb{P}(S_{\text{ogp}}^c)$.

Lemma 5.3. *Suppose $y^{(0)}, \dots, y^{(k)} \in \{\mathbf{T}, \mathbf{F}\}^n$ is a sequence of assignments and $0 \leq t_0 \leq t_1 \leq \dots \leq t_k \leq T$. Then,*

$$\frac{1}{n} \log \mathbb{P} \left[y^{(\ell)} \text{ satisfies } \Phi^{(t_\ell)} \text{ for all } 0 \leq \ell \leq k \right] \leq -\kappa \frac{\log k}{k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right| + o(1).$$

Proof. For $1 \leq i \leq m$, the collections of clauses $\{\Phi_i^{(t)} : 0 \leq t \leq T\}$ are mutually independent. So,

$$\mathbb{P} \left[y^{(\ell)} \text{ satisfies } \Phi^{(t_\ell)} \text{ for all } 0 \leq \ell \leq k \right] = \prod_{i=1}^m \mathbb{P} \left[y^{(\ell)} \text{ satisfies } \Phi_i^{(t_\ell)} \text{ for all } 0 \leq \ell \leq k \right].$$

We say a clause index $i \in [m]$ is *interrupted* if for some $0 \leq \ell \leq k$, t_ℓ satisfies $1 \leq \sigma(t_\ell) - (i-1)k \leq k-1$, where $\sigma(\cdot)$ is defined in Definition 4.2. Informally, i is interrupted if there is some ℓ such that $\Phi^{(t_\ell)}$ is partway

through resampling the i th clause. Note that each t_ℓ interrupts at most one clause, so there are at most $k+1$ interrupted clause indices. Because so few clause indices are interrupted, it does not hurt our analysis to throw them out. We clearly have

$$\mathbb{P}\left[y^{(\ell)} \text{ satisfies } \Phi^{(t_\ell)} \text{ for all } 0 \leq \ell \leq k\right] \leq \prod_{\substack{i \in [m] \\ \text{not interrupted}}} \mathbb{P}\left[y^{(\ell)} \text{ satisfies } \Phi_i^{(t_\ell)} \text{ for all } 0 \leq \ell \leq k\right]. \quad (5.2)$$

We now fix a single non-interrupted index $i \in [m]$ and analyze the last probability. We exploit the following stochastic property of non-interrupted clauses: if i is not interrupted, then the clauses $\Phi_i^{(t_0)}, \Phi_i^{(t_1)}, \dots, \Phi_i^{(t_k)}$ can be partitioned into equivalence classes, such that all clauses in the same equivalence class are identical and all clauses in different equivalence classes are mutually independent. Formally, for some $1 \leq r \leq k+1$, there is a surjective map $\tau : \{0, \dots, k\} \rightarrow [r]$ (dependent only on the indices $t_0, \dots, t_k$ and i) such that for i.i.d. clauses $C_1, \dots, C_r \sim \Phi_k(n, 1)$,

$$\left(\Phi_i^{(t_0)}, \Phi_i^{(t_1)}, \dots, \Phi_i^{(t_k)}\right) =_d (C_{\tau(0)}, C_{\tau(1)}, \dots, C_{\tau(k)}).$$

For $1 \leq s \leq r$, let $B_s = \tau^{-1}(s)$ be the set of $\ell \in \{0, \dots, k\}$ such that $\Phi_i^{(t_\ell)}$ corresponds to C_s . Thus $B_1, \dots, B_r$ partition $\{0, \dots, k\}$. Now,

$$\mathbb{P}\left[y^{(\ell)} \text{ satisfies } \Phi_i^{(t_\ell)} \text{ for all } 0 \leq \ell \leq k\right] = \prod_{s=1}^r \mathbb{P}\left[y^{(\ell)} \text{ satisfies } C_s \text{ for all } \ell \in B_s\right]. \quad (5.3)$$

Let $I \in [n]^k$ be the indices of the k variables sampled by C_s , so $I \sim \text{unif}([n]^k)$. Given I , there are $|\{y^{(\ell)}[I] : \ell \in B_s\}|$ ways to assign polarities to these k variables such that for some $\ell \in B_s$, $y^{(\ell)}$ does not satisfy C_s . Thus, conditioned on I , the probability that $y^{(\ell)}$ satisfies C_s for all $\ell \in B_s$ is $1 - 2^{-k} |\{y^{(\ell)}[I] : \ell \in B_s\}|$. It follows that

$$\begin{aligned} \mathbb{P}\left[y^{(\ell)} \text{ satisfies } C_s \text{ for all } \ell \in B_s\right] &= 1 - 2^{-k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \{y^{(\ell)}[I] : \ell \in B_s\} \right| \\ &\leq \exp \left(-2^{-k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \{y^{(\ell)}[I] : \ell \in B_s\} \right| \right). \end{aligned}$$

So, using (5.3) and recalling that $B_1, \dots, B_r$ partition $\{0, \dots, k\}$, we have

$$\begin{aligned} \mathbb{P}\left[y^{(\ell)} \text{ satisfies } \Phi_i^{(t_\ell)} \text{ for all } 0 \leq \ell \leq k\right] &\leq \exp \left(-2^{-k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \sum_{s=1}^r \left| \{y^{(\ell)}[I] : \ell \in B_s\} \right| \right) \\ &\leq \exp \left(-2^{-k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \{y^{(\ell)}[I] : 0 \leq \ell \leq k\} \right| \right). \end{aligned}$$

Next, we substitute into (5.2). There are at most $k+1$ interrupted clauses, and thus at least $m - (k+1) \geq n\alpha - k - 2$ non-interrupted clauses. So,

$$\mathbb{P}\left[y^{(\ell)} \text{ satisfies } \Phi^{(t_\ell)} \text{ for all } 0 \leq \ell \leq k\right] \leq \exp \left(-(n\alpha - k - 2) 2^{-k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \{y^{(\ell)}[I] : 0 \leq \ell \leq k\} \right| \right).$$

Thus,

$$\frac{1}{n} \log \mathbb{P}\left[y^{(\ell)} \text{ satisfies } \Phi^{(t_\ell)} \text{ for all } 0 \leq \ell \leq k\right] \leq -\alpha 2^{-k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \{y^{(\ell)}[I] : 0 \leq \ell \leq k\} \right| + o(1).$$

The result follows from $\alpha = \kappa 2^k \log k/k$. $\square$

Proof of Lemma 5.1. By Markov's inequality, $\mathbb{P}(S_{\text{ogp}}^c)$ is upper bounded by the expected number of $0 \leq t_0 \leq t_1 \leq \dots \leq t_k \leq T$ and $(y^{(0)}, \dots, y^{(k)})$ satisfying conditions (OGP-A) and (OGP-B) of S_{ogp} . There are at

most $(T+1)^{k+1}$ choices of $(t_0, \dots, t_k)$, and (by Fact 4.5(a)) $|P| \leq n^{2^k}$. By linearity of expectation,

$$\mathbb{P}(S_{\text{ogp}}^c) \leq (T+1)^{k+1} n^{2^k} \max_{\substack{0 \leq t_0 \leq \dots \leq t_k \leq T \\ \pi \in P}} \mathbb{E} \left[\# \left(\begin{array}{l} (y^{(0)}, \dots, y^{(k)}) \in \{\mathbf{T}, \mathbf{F}\}^{n \times (k+1)} : \\ y^{(\ell)} \text{ satisfies } \Phi^{(t_\ell)} \text{ for all } 0 \leq \ell \leq k \\ \text{and } \pi(y^{(0)}, \dots, y^{(k)}) = \pi \end{array} \right) \right].$$

Let πn be the scalar product of π , treated as a vector, by n . There are $2^n \binom{n}{\pi n}$ sequences of assignments $(y^{(0)}, \dots, y^{(k)})$ with $\pi(y^{(0)}, \dots, y^{(k)}) = \pi$: 2^n ways to choose $y^{(0)}$, and then $\binom{n}{\pi n}$ ways to assign the positions $[n]$ to the partitions of $\{0, \dots, k\}$. Over all of these sequences of assignments, the probability of the event that $y^{(\ell)}$ satisfies $\Phi^{(t_\ell)}$ for all $0 \leq \ell \leq k$ is uniformly upper bounded by Lemma 5.3. By linearity of expectation, the last expectation is upper bounded by

$$2^n \binom{n}{\pi n} \exp \left(-n\kappa \frac{\log k}{k} \mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \{y^{(\ell)}[I] : 0 \leq \ell \leq k\} \right| + o(n) \right).$$

Because $\binom{n}{\pi n} = \exp(nH(\pi) + o(n))$, the result follows. $\square$

Remark 5.4. The step in the proof of Lemma 5.3 where we lower bound $\sum_{s=1}^r |\{y^{(\ell)}[I] : \ell \in B_s\}|$ by $|\{y^{(\ell)}[I] : 0 \leq \ell \leq k\}|$ is tight when $t_0, \dots, t_k$ are all equal, because in this case $r = 1$ and $B_1 = \{0, 1, \dots, k\}$. Thus the exponential rate of $\mathbb{P}(S_{\text{ogp}}^c)$ is dominated by the case when the t_i are equal. In other words, $\mathbb{P}(S_{\text{ogp}}^c)$ has the same exponential rate as if, in the definition of S_{ogp} , we required all the $y^{(\ell)}$ to be satisfying assignments to the same $\Phi^{(t)}$. This shows the power of the “ensemble” part of the ensemble multi-OGP: for no cost in the exponential rate, we can generalize the forbidden structure to an ensemble, which we create by running a smooth algorithm on all problem instances in an interpolation path. All ensemble (multi-)OGPs in the literature share this property, see [43, 71].

5.3 Lower Bounding the Energy Term

In this section, we will prove Proposition 5.2. Let $y^{(0)}, \dots, y^{(k)}$ and $\beta_1, \dots, \beta_k$ be as in Proposition 5.2. Without loss of generality, we can set $y^{(0)} = \mathbf{T}^n$.

To analyze the expectation in Proposition 5.2, we introduce the following probabilistic quantities. For $0 \leq \ell \leq k$ and $\sigma \in \{\mathbf{T}, \mathbf{F}\}^k$, define

$$E_\ell(\sigma) = \left\{ I \in [n]^k : y^{(\ell')}[I] = \sigma \text{ for some } 0 \leq \ell' \leq \ell \right\} \quad \text{and} \quad p_\ell(\sigma) = \mathbb{P}_{I \sim \text{unif}([n]^k)}(E_\ell(\sigma)).$$

In other words, $E_\ell(\sigma)$ is the event that σ appears in the set $\{y^{(\ell')}[I] : 0 \leq \ell' \leq \ell\}$, and $p_\ell(\sigma)$ is the probability of this event. The probabilities $p_k(\sigma)$ will be relevant to our analysis by the following identity (5.4), while the probabilities $p_\ell(\sigma)$ for $\ell < k$ will arise in our inductive analysis below, where we lower bound $p_k(\sigma)$ by peeling off one of $y^{(1)}, \dots, y^{(k)}$ at a time. We have that

$$\mathbb{E}_{I \sim \text{unif}([n]^k)} \left| \{y^{(\ell)}[I] : 0 \leq \ell \leq k\} \right| = \mathbb{E}_{I \sim \text{unif}([n]^k)} \left[\sum_{\sigma \in \{\mathbf{T}, \mathbf{F}\}^k} \mathbb{1}\{I \in E_k(\sigma)\} \right] = \sum_{\sigma \in \{\mathbf{T}, \mathbf{F}\}^k} p_k(\sigma). \quad (5.4)$$

To prove Proposition 5.2, we will need to lower bound the right-hand side of (5.4). This task will require several definitions; to motivate these definitions, let us first outline our technique for deriving this lower bound. The first idea of our technique is a conditional expansion. We can reveal the k bit strings in the tuple $(y^{(1)}[I], \dots, y^{(k)}[I])$ one by one; conditioned on its predecessors $y^{(1)}[I], \dots, y^{(\ell-1)}[I]$, the distribution of $y^{(\ell)}[I]$ can be described in terms of the conditional overlap profile $\pi(y^{(\ell)} | y^{(0)}, \dots, y^{(\ell-1)})$ (recall here that we set $y^{(0)} = \mathbf{T}^n$). Then, $1 - p_k(\sigma)$, which is the probability that σ does not appear in the set $\{y^{(\ell)}[I] : 0 \leq \ell \leq k\}$, can be expanded as a product of k factors: the ℓ th factor is the (conditional) probability that the revealed value of $y^{(\ell)}$ does not equal σ . Thus the ℓ th factor of this product can be thought of as the contribution of $y^{(\ell)}$.

Our second idea is to estimate this product by a *sum*, whose ℓ th summand is the contribution of $y^{(\ell)}$. The purpose of this estimation is to decouple the contributions of the $y^{(\ell)}$, so that we can analyze the overall

contribution of $y^{(\ell)}$ by summing over $\sigma \in \{\mathbf{T}, \mathbf{F}\}^k$. We achieve this by truncating the factors in the product at $1 - \frac{1}{k \log k}$; any factor smaller than this gets rounded up to 1. Then we can, by an inductive argument, separate off the contributions of $y^{(1)}, \dots, y^{(k)}$ to $p_k(\sigma)$ one by one. We choose $\frac{1}{k \log k}$ to be slightly smaller than $\frac{1}{k}$, so that the resulting estimation by a sum results in $1 - o_k(1)$ multiplicative error. Propositions 5.5 and 5.6 below carry out this technique. Finally, our third idea is to collect the (now additive) contribution of each $y^{(\ell)}$ over all $\sigma \in \{\mathbf{T}, \mathbf{F}\}^k$. Miraculously, we can reformulate this contribution as a probability of a sum of k i.i.d. random variables, which can be controlled by a Chernoff bound. This step is carried out in Proposition 5.7.

Formally, for $0 \leq \ell \leq k$ and $i \in [n]$, let $y_i^{(\leq \ell)} = (y_i^{(1)}, \dots, y_i^{(\ell)})$. Similarly, for $I \in [n]^k$, let $y^{(\leq \ell)}[I] = (y^{(1)}[I], \dots, y^{(\ell)}[I])$. Because $y^{(0)} = \mathbf{T}^n$, the overlap profile π determines the distribution of $y_i^{(\leq k)}$, where $i \sim \text{unif}([n])$. Namely, for $\tau \in \{\mathbf{T}, \mathbf{F}\}^k$,

$$\mathbb{P}_{i \sim \text{unif}([n])} \left[y_i^{(\leq k)} = \tau \right] = \pi_{S \cup \{0\}, T}$$

where $S = \{\ell \in [k] : \tau_\ell = \mathbf{T}\}$ and $T = \{\ell \in [k] : \tau_\ell = \mathbf{F}\}$. Moreover, the distribution of $y^{(\leq k)}[I]$, where $I \sim \text{unif}([n]^k)$, is the product of k i.i.d. copies of this distribution. For $1 \leq \ell \leq k$, $b \in \{\mathbf{T}, \mathbf{F}\}$, and $\tau \in \{\mathbf{T}, \mathbf{F}\}^{\ell-1}$, define

$$\phi_\ell(b|\tau) = \mathbb{P}_{i \sim \text{unif}([n])} \left[y_i^{(\ell)} = b | y_i^{(\leq \ell-1)} = \tau \right].$$

The probabilities in the aforementioned conditional expansion are products of conditional probabilities $\phi_\ell(b|\tau)$. Namely, the probability that $y^{(\ell)}[I] \neq \sigma$ given $y^{(\leq \ell-1)}[I]$ is $1 - \prod_{r=1}^k \phi_\ell(\sigma_r | y_{I_r}^{(\leq \ell-1)})$.

For $1 \leq \ell \leq k$, $\sigma \in \{\mathbf{T}, \mathbf{F}\}^k$ and $I \in [n]^k$, further define

$$Q_\ell(\sigma, I) = \prod_{r=1}^k \phi_\ell(\sigma_r | y_{I_r}^{(\leq \ell-1)}) \mathbb{1} \left\{ \prod_{r=1}^k \phi_\ell(\sigma_r | y_{I_r}^{(\leq \ell-1)}) \leq \frac{1}{k \log k} \right\} \quad \text{and} \quad q_\ell(\sigma) = \mathbb{E}_{I \sim \text{unif}([n]^k)} [Q_\ell(\sigma, I)].$$

Thus, $1 - Q_\ell(\sigma, I)$ is a term in the conditional expansion, truncated at $1 - \frac{1}{k \log k}$ in the aforementioned sense, and $q_\ell(\sigma)$ is its expectation.

For each $\sigma \in \{\mathbf{T}, \mathbf{F}\}^k$, the following two propositions lower bound $p_k(\sigma)$ in terms of $q_1(\sigma), \dots, q_k(\sigma)$ by peeling off one of $y^{(1)}, \dots, y^{(k)}$ at a time.

Proposition 5.5. *For each $\sigma \in \{\mathbf{T}, \mathbf{F}\}^k$ and $1 \leq \ell \leq k$, we have that*

$$p_\ell(\sigma) \geq \left(1 - \frac{1}{k \log k} \right) p_{\ell-1}(\sigma) + q_\ell(\sigma).$$

Proof. Note that

$$\begin{aligned} 1 - p_\ell(\sigma) &= \mathbb{P}_{I \sim \text{unif}([n]^k)} \left[y^{(\ell')}[I] \neq \sigma \text{ for all } 0 \leq \ell' \leq \ell \right] \\ &= \mathbb{E}_{I \sim \text{unif}([n]^k)} \left[\mathbb{1} \left\{ y^{(\ell')}[I] \neq \sigma \text{ for all } 0 \leq \ell' \leq \ell - 1 \right\} \left(1 - \prod_{r=1}^k \phi_\ell(\sigma_r | y_{I_r}^{(\leq \ell-1)}) \right) \right]. \end{aligned}$$

Here, we use that the event inside the indicator is $y^{(\leq \ell-1)}[I]$ -measurable, and conditioned on $y^{(\leq \ell-1)}[I]$ the probability that $y^{(\ell)}[I] = \sigma$ is $\prod_{r=1}^k \phi_\ell(\sigma_r | y_{I_r}^{(\leq \ell-1)})$. Moreover, we have $\prod_{r=1}^k \phi_\ell(\sigma_r | y_{I_r}^{(\leq \ell-1)}) \geq Q_\ell(\sigma, I)$ by definition. So,

$$\begin{aligned} 1 - p_\ell(\sigma) &\leq \mathbb{E}_{I \sim \text{unif}([n]^k)} \left[(1 - \mathbb{1} \{I \in E_{\ell-1}(\sigma)\}) (1 - Q_\ell(\sigma, I)) \right] \\ &\leq \mathbb{E}_{I \sim \text{unif}([n]^k)} \left[1 - \left(1 - \frac{1}{k \log k} \right) \mathbb{1} \{I \in E_{\ell-1}(\sigma)\} - Q_\ell(\sigma, I) \right] \\ &= 1 - \left(1 - \frac{1}{k \log k} \right) p_{\ell-1}(\sigma) - q_\ell(\sigma). \end{aligned}$$

The second-last line uses the fact that $Q_\ell(\sigma, I) \leq \frac{1}{k \log k}$ almost surely, and the last line uses the definitions of $p_{\ell-1}(\sigma)$ and $q_\ell(\sigma)$. Rearranging yields the desired bound. $\square$

Proposition 5.6. For each $\sigma \in \{\mathbf{T}, \mathbf{F}\}^k$, we have that

$$p_k(\sigma) \geq \left(1 - \frac{1}{\log k}\right) \sum_{\ell=1}^k q_\ell(\sigma).$$

Proof. By iterating Proposition 5.5, we get

$$p_k(\sigma) \geq \left(1 - \frac{1}{k \log k}\right)^k p_0(\sigma) + \sum_{\ell=1}^k \left(1 - \frac{1}{k \log k}\right)^{k-\ell} q_\ell(\sigma) \geq \left(1 - \frac{1}{k \log k}\right)^k \sum_{\ell=1}^k q_\ell(\sigma).$$

The result follows from the bound $\left(1 - \frac{1}{k \log k}\right)^k \geq 1 - \frac{1}{\log k}$, by Bernoulli's inequality. $\square$

Equation (5.4) and Proposition 5.6 leave the task of lower bounding $\sum_{\sigma \in \{\mathbf{T}, \mathbf{F}\}^k} \sum_{\ell=1}^k q_\ell(\sigma)$. This is achieved by the following proposition, which reinterprets $\sum_{\sigma \in \{\mathbf{T}, \mathbf{F}\}^k} q_\ell(\sigma)$, the total contribution of $y^{(\ell)}$, as a probability.

Proposition 5.7. For each $1 \leq \ell \leq k$, we have that

$$\sum_{\sigma \in \{\mathbf{T}, \mathbf{F}\}^k} q_\ell(\sigma) \geq 1 - \beta_\ell e^{-(\beta_\ell - 1)} - o_k(1).$$

Proof. Using the definition of $q_\ell(\sigma)$, we have

$$\begin{aligned} \sum_{\sigma \in \{\mathbf{T}, \mathbf{F}\}^k} q_\ell(\sigma) &= \mathbb{E}_{I \sim \text{unif}([n]^k)} \left[\sum_{\sigma \in \{\mathbf{T}, \mathbf{F}\}^k} \prod_{r=1}^k \phi_\ell(\sigma_r | y_{I_r}^{(\leq \ell-1)}) \mathbb{1} \left\{ \prod_{r=1}^k \phi_\ell(\sigma_r | y_{I_r}^{(\leq \ell-1)}) \leq \frac{1}{k \log k} \right\} \right] \\ &= \mathbb{E}_{I \sim \text{unif}([n]^k)} \left[\sum_{\sigma \in \{\mathbf{T}, \mathbf{F}\}^k} \prod_{r=1}^k \phi_\ell(\sigma_r | y_{I_r}^{(\leq \ell-1)}) \mathbb{1} \left\{ -\sum_{r=1}^k \log \phi_\ell(\sigma_r | y_{I_r}^{(\leq \ell-1)}) \geq \log k + \log \log k \right\} \right]. \end{aligned}$$

This quantity is the success probability of the following experiment. Sample positive random variables $u_1, \dots, u_k$ by the following procedure, repeated independently for each $r \in [k]$. Sample $i \in \text{unif}([n])$; this determines the value of $y_i^{(\leq \ell-1)}$. Then, sample $b \in \{\mathbf{T}, \mathbf{F}\}$ from the measure $\phi_\ell(\cdot | y_i^{(\leq \ell-1)})$. Finally, set $u_r = -\log \phi_\ell(b | y_{I_r}^{(\leq \ell-1)})$. The experiment succeeds if $\sum_{r=1}^k u_r \geq \log k + \log \log k$.

For $r \in [k]$, let $v_r = \min(u_r, \log k)$. Informally, v_r is a proxy for u_r with an almost sure upper bound, which allows us to control the experiment's failure probability by a Chernoff bound. This failure probability is bounded by

$$\mathbb{P} \left[\sum_{r=1}^k u_r < \log k + \log \log k \right] \leq \mathbb{P} \left[\sum_{r=1}^k v_r < \log k + \log \log k \right] = \mathbb{P} \left[\sum_{r=1}^k \frac{v_r}{\log k} < 1 + \frac{\log \log k}{\log k} \right].$$

Note that the $\frac{v_r}{\log k}$ are i.i.d. random variables in $[0, 1]$ almost surely. To bound this last probability by a Chernoff bound, we will lower bound $\mathbb{E}[v_r]$. By the definition of ϕ_ℓ ,

$$\begin{aligned} \mathbb{E}[u_r] &= \mathbb{E}_{i \sim \text{unif}([n])} \mathbb{E}_{b \sim \phi_\ell(\cdot | y_i^{(\leq \ell-1)})} \left[-\log \phi_\ell(b | y_i^{(\leq \ell-1)}) \right] \\ &= H(\pi(y^{(\ell)} | y^{(0)}, \dots, y^{(\ell-1)})) = \beta_\ell \frac{\log k}{k}. \end{aligned}$$

Moreover,

$$\begin{aligned} \mathbb{E}[u_r - v_r] &= \mathbb{E}[(u_r - \log k) \mathbb{1}\{u_r \geq \log k\}] \\ &= \mathbb{E}_{i \sim \text{unif}([n])} \left[\sum_{b \in \{\mathbf{T}, \mathbf{F}\}} \phi_\ell(b | y_i^{(\leq \ell-1)}) \log \frac{1}{k \phi_\ell(b | y_i^{(\leq \ell-1)})} \mathbb{1} \left\{ \phi_\ell(b | y_i^{(\leq \ell-1)}) \leq \frac{1}{k} \right\} \right] \end{aligned}$$

For each i , the quantity inside the last expectation is nonzero for at most one $b \in \{\mathbf{T}, \mathbf{F}\}$ (for $k \geq 3$). Moreover, on the interval $[0, \frac{1}{k}]$, the function $x \mapsto x \log \frac{1}{kx}$ has maximum value $\frac{1}{ek}$, attained at $x = \frac{1}{ek}$. Thus, $\mathbb{E}[u_r - v_r] \leq \frac{1}{ek}$. It follows that $\mathbb{E}[v_r] \geq \beta_\ell \frac{\log k}{k} - \frac{1}{ek}$. So,

$$\mathbb{E} \left[\sum_{r=1}^k \frac{v_r}{\log k} \right] \geq \beta_\ell - \frac{1}{e \log k}.$$

Furthermore, $(1 + \frac{\log \log k}{\log k}) / (\beta_\ell - \frac{1}{e \log k}) = \frac{1}{\beta_\ell} + o_k(1)$. So, by a Chernoff bound,

$$\mathbb{P} \left[\sum_{r=1}^k \frac{v_r}{\log k} < 1 + \frac{\log \log k}{\log k} \right] \leq \left(\frac{e^{-(1 - \frac{1}{\beta_\ell} - o_k(1))}}{\left(\frac{1}{\beta_\ell} + o_k(1) \right)^{\frac{1}{\beta_\ell} + o_k(1)}} \right)^{\beta_\ell - o_k(1)} = \beta_\ell e^{-(\beta_\ell - 1)} + o_k(1).$$

Hence,

$$\mathbb{P} \left[\sum_{r=1}^k u_r \geq \log k + \log \log k \right] \geq 1 - \beta_\ell e^{-(\beta_\ell - 1)} - o_k(1),$$

as desired. $\square$

We can now combine these propositions to prove Proposition 5.2.

Proof of Proposition 5.2. By combining (5.4), Proposition 5.6, and Proposition 5.7, we have

$$\begin{aligned} \mathbb{E}_{I \sim \text{unif}(\{[n]^k\})} \left| \left\{ y^{(\ell)}[I] : 0 \leq \ell \leq k \right\} \right| &\geq \left(1 - \frac{1}{\log k} \right) \sum_{\ell=1}^k \sum_{\sigma \in \{\mathbf{T}, \mathbf{F}\}^k} q_\ell(\sigma) \\ &\geq (1 - o_k(1)) \sum_{\ell=1}^k \left(1 - \beta_\ell e^{-(\beta_\ell - 1)} \right). \end{aligned}$$

$\square$

6 Stability of Low Degree Polynomials

In this section, we will prove Proposition 4.7(a), which lower bounds the probability that $x^{(t)}$ satisfies $\Phi^{(t)}$ for all $0 \leq t \leq T$ and the sequence $x^{(t)}$ has no large jumps in Hamming distance.

6.1 An Upper Bound on the Rate of Bad Steps

We begin by defining the notion of c -badness, which will be crucial to our proof. Informally, (Φ, Φ') is c -bad if the output of f has a large jump between inputs Φ and Φ' . Recall that $N = m \cdot k \cdot 2n$, and each $\Phi \in \Omega_k(n, m)$ can be identified with a vector of indicators in $\{0, 1\}^N$, which is the input of a low degree polynomial.

Definition 6.1 (c -badness). Let $c > 0$ and let $f : \mathbb{R}^N \rightarrow \mathbb{R}^n$ be a deterministic degree- D polynomial. A pair of formulas $(\Phi, \Phi') \in \Omega_k(n, m)^2$ is c -bad (with respect to f) if $\|f(\Phi) - f(\Phi')\|_2^2 > c \mathbb{E}_{\Phi \sim \Phi_k(n, m)} \|f(\Phi)\|_2^2$.

Recall the interpolation path $\Phi^{(0)}, \Phi^{(1)}, \dots, \Phi^{(T)}$ defined in Definition 4.2. We will prove Proposition 4.7(a) via the following proposition, which controls the probability that the output of f does not have a large jump between any pair of consecutive assignments in the interpolation path.

Proposition 6.2. Let $f : \mathbb{R}^N \rightarrow \mathbb{R}^n$ be a deterministic degree- D polynomial. With probability at least $(2n)^{-4Dk/c}$, $(\Phi^{(t-1)}, \Phi^{(t)})$ is not c -bad with respect to f for any $1 \leq t \leq T$.

We will prove this proposition in Section 6.3. The objective of this subsection is to prove Proposition 6.3 below, which upper bounds the fraction of possible steps that can be bad. To this end, for $1 \leq j \leq km$, define $\Phi_k(n, m; j)$ as the measure of a sample $(\Phi, \Phi') \in \Omega_k(n, m)^2$ obtained by sampling $\Phi \sim \Phi_k(n, m)$, and then obtaining Φ' from Φ by resampling the j th lexicographic literal $L(\Phi, j)$ from $\text{unif}(\mathcal{L} \setminus \{L(\Phi, j)\})$. Define

$$\lambda_j = \mathbb{P}_{(\Phi, \Phi') \sim \Phi_k(n, m; j)} ((\Phi, \Phi') \text{ is } c\text{-bad with respect to } f).$$

This is the fraction of pairs of formulas in $\Omega_k(n, m)$, differing in exactly the j th lexicographic literal, that are c -bad with respect to f .

Proposition 6.3. *If f is a deterministic degree- D polynomial, then $\sum_{j=1}^{km} \lambda_j \leq \frac{4D}{c}$.*

We recall the following orthogonal decomposition property of functions on product measures, which can be thought of as a generalization of Fourier analysis on the Boolean cube. We will give brief self-contained proofs of the relevant facts; a full discussion can be found in [67, Chapter 8.3]. Let π be a probability measure on an arbitrary space $\mathcal{X}$, and let J be a positive integer. Let $X = (X_1, \dots, X_J) \in \mathcal{X}^J$. For $j \in [J]$, define the operators D_j and E_j as follows. For any function $g : \mathcal{X}^J \rightarrow \mathbb{R}$, $E_j g$ is the function satisfying

$$E_j g(X) = \mathbb{E}_{X_j \sim \pi} g(X),$$

where in the right-hand side the coordinate X_j of X is resampled from π . Let $D_j g = g - E_j g$. Note that the operators D_j, E_j commute (including across multiple $j \in [J]$). For $S \subseteq [J]$, define the functions

$$\hat{g}_S = \prod_{j \in S} D_j \prod_{j \in [J] \setminus S} E_j g.$$

Note that $g = \sum_{S \subseteq [J]} \hat{g}_S$. Moreover, $\hat{g}_S$ depends only on the inputs $\{X_j : j \in S\}$. For any j ,

$$\mathbb{E}_{X_j \sim \pi} g(X)^2 = \mathbb{E}_{X_j \sim \pi} [(D_j g)(X)^2] + (E_j g)(X)^2,$$

and so by induction

$$\mathbb{E}_{X \sim \pi^{\otimes J}} g(X)^2 = \sum_{S \subseteq [J]} \mathbb{E}_{X \sim \pi^{\otimes J}} \hat{g}_S(X)^2.$$

For $j \in [J]$, define $\text{Var}_j g(X) = \mathbb{E}_{X_j \sim \pi} [(D_j g)(X)^2]$. We begin with the following inequality, which can be considered a converse to the Efron-Stein inequality.

Lemma 6.4. *Suppose a function $g : \mathcal{X}^J \rightarrow \mathbb{R}$ can be written in the form $g(X) = \sum_{i=1}^I g_i(X)$, where each $g_i(X)$ depends on at most D coordinates of X . Then,*

$$D \mathbb{E}_{X \sim \pi^{\otimes J}} \text{Var}_j g(X) \geq \sum_{j=1}^J \mathbb{E}_{X \sim \pi^{\otimes J}} \text{Var}_j g(X).$$

Proof. By the orthogonal expansion above, we have

$$\mathbb{E}_{X \sim \pi^{\otimes J}} \text{Var}_j g(X) = \sum_{\substack{S \subseteq [J] \\ S \neq \emptyset}} \mathbb{E}_{X \sim \pi^{\otimes J}} \hat{g}_S(X)^2 \quad \text{and} \quad \mathbb{E}_{X \sim \pi^{\otimes J}} \text{Var}_j g(X) = \sum_{\substack{S \subseteq [J] \\ S \ni j}} \mathbb{E}_{X \sim \pi^{\otimes J}} \hat{g}_S(X)^2.$$

We claim that for all $S \subseteq [J]$ with $|S| > D$, we have $\hat{g}_S \equiv 0$. For each $i \in [I]$, we have $\prod_{j \in S} D_j g_i \equiv 0$, because S contains at least one j such that $g_i(X)$ does not depend on X_j . Thus, $\prod_{j \in S} D_j g \equiv 0$, and so $\hat{g}_S \equiv 0$, as desired. Hence,

$$\sum_{j=1}^J \mathbb{E}_{X \sim \pi^{\otimes J}} \text{Var}_j g(X) = \sum_{S \subseteq [J]} |S| \mathbb{E}_{X \sim \pi^{\otimes J}} \hat{g}_S(X)^2 \leq D \mathbb{E}_{X \sim \pi^{\otimes J}} \text{Var}_j g(X).$$

□

Proof of Proposition 6.3. Note that $\Phi_k(n, m)$ is composed of km i.i.d. literals, and thus can be thought of as the product measure $\text{unif}(\mathcal{L})^{\otimes km}$. By slight abuse of notation, for $1 \leq j \leq km$, we can define D_j and E_j as the above operators with respect to the j th lexicographic literal of Φ .

For $1 \leq \ell \leq n$, let f_ℓ denote the ℓ th component of f . By Markov's inequality and the inequality $(a - b)^2 \leq 2a^2 + 2b^2$, we have

$$\begin{aligned} \sum_{j=1}^{km} \lambda_j &\leq \sum_{j=1}^{km} \frac{\mathbb{E}_{(\Phi, \Phi') \sim \Phi_k(n, m; j)} \|f(\Phi) - f(\Phi')\|_2^2}{c \mathbb{E}_{\Phi \sim \Phi_k(n, m)} \|f(\Phi)\|_2^2} \\ &= \frac{\sum_{\ell=1}^n \sum_{j=1}^{km} \mathbb{E}_{(\Phi, \Phi') \sim \Phi_k(n, m; j)} ((D_j f_\ell)(\Phi) - (D_j f_\ell)(\Phi'))^2}{c \sum_{\ell=1}^n \mathbb{E}_{\Phi \sim \Phi_k(n, m)} f_\ell(\Phi)^2} \\ &\leq \frac{2 \sum_{\ell=1}^n \sum_{j=1}^{km} \mathbb{E}_{(\Phi, \Phi') \sim \Phi_k(n, m; j)} ((D_j f_\ell)(\Phi))^2 + ((D_j f_\ell)(\Phi'))^2}{c \sum_{\ell=1}^n \mathbb{E}_{\Phi \sim \Phi_k(n, m)} f_\ell(\Phi)^2} \\ &= \frac{4 \sum_{\ell=1}^n \sum_{j=1}^{km} \mathbb{E}_{\Phi \sim \Phi_k(n, m)} \text{Var}_j f_\ell(\Phi)}{c \sum_{\ell=1}^n \mathbb{E}_{\Phi \sim \Phi_k(n, m)} f_\ell(\Phi)^2}. \end{aligned}$$

Now, each f_ℓ is a degree- D polynomial in the indicators $\Phi_{i,j,s}$ that $\Phi_{i,j}$ is the s th literal in $\mathcal{L}$. So, each monomial of each f_ℓ depends on at most D literals of Φ . By Lemma 6.4,

$$\sum_{j=1}^{km} \mathbb{E}_{\Phi \sim \Phi_k(n, m)} \text{Var}_j f_\ell(\Phi) \leq D \text{Var}_{\Phi \sim \Phi_k(n, m)} f_\ell(\Phi) \leq D \mathbb{E}_{\Phi \sim \Phi_k(n, m)} f_\ell(\Phi)^2.$$

So, $\sum_{j=1}^{km} \lambda_j \leq \frac{4D}{c}$. □

6.2 Bounding the Probability of no Bad Step

Proposition 6.3 bounds the combined rate of c -bad steps. To derive Proposition 6.2, we must translate this bound on the rate of c -bad steps to a bound on the probability that interpolation path never takes a c -bad step. To make the ideas in our argument more clear, we abstract to the following graph theoretic problem, which is interesting in its own right.

Let Σ be a set of symbols and J, T be positive integers. Let G be a graph on Σ^J , where two nodes are adjacent if their Hamming distance is exactly 1. Each edge has a *direction* $j \in [J]$, the index on which its endpoints disagree. Let an arbitrary subset of edges be *bad*; for adjacent vertices v, w , let $B(v, w)$ denote the event that the edge (v, w) is bad. For $j \in [J]$, let λ_j denote the fraction of edges in direction j that are bad. Equivalently, $\lambda_j = \mathbb{P}(B(v, w))$, where $v \sim \text{unif}(G)$ and w is obtained from v by resampling w_j from $\text{unif}(\Sigma \setminus \{v_j\})$.

Let $\sigma : [T] \rightarrow [J]$ be an arbitrary map. Consider the (lazy) random walk $v^{(0)}, v^{(1)}, \dots, v^{(T)}$ such that $v^{(0)} \sim \text{unif}(G)$ and for $1 \leq t \leq T$, $v^{(t)}$ is obtained from $v^{(t-1)}$ by resampling $v_{\sigma(t)}^{(t)}$ from $\text{unif}(\Sigma)$.

Lemma 6.5. *With probability at least $|\Sigma|^{-\sum_{t=1}^T \lambda_{\sigma(t)}}$, no step of the random walk $v^{(0)}, v^{(1)}, \dots, v^{(T)}$ traverses a bad edge.*

Note that at each step, the random walk either traverses an edge or does not move; we say that the steps that do not move do not traverse a bad edge. The lemma is sharp, for example, when all the λ_j are 0 or 1: in this case, the random walk does not traverse a bad edge if it does not move at all times t with $\lambda_{\sigma(t)} = 1$.

Proof. For $v \in \Sigma^J$, let $q(v)$ be the probability that the random walk $v^{(0)}, v^{(1)}, \dots, v^{(T)}$ does not traverse a bad edge, starting from $v^{(0)} = v$. We will prove by induction on T that

$$\mathbb{E}_{v \sim \text{unif}(G)} \log q(v) \geq -\log |\Sigma| \cdot \sum_{t=1}^T \lambda_{\sigma(t)}.$$

The lemma then follows from Jensen's inequality, because $\log \mathbb{E} q(v) \geq \mathbb{E} \log q(v)$.

The base case of the claim, $T = 0$, follows trivially. For the inductive step, let $\tilde{q}(v)$ be the probability that the random walk $v^{(1)}, v^{(2)}, \dots, v^{(T)}$ does not traverse a bad edge, starting from $v^{(1)} = v$. Let $j = \sigma(1)$. Let $v_{\sim j} \in \Sigma^{J-1}$ denote an element of Σ^J with the j th coordinate left blank. For $s \in \Sigma$, let $v_{\sim j}[s] \in \Sigma^J$ denote $v_{\sim j}$ with the j th coordinate set to s .

For now, fix some $v_{\sim j} \in \Sigma^{J-1}$. For $s \in \Sigma$, we have that

$$\begin{aligned} q(v_{\sim j}[s]) &= \sum_{s' \in \Sigma} \frac{1}{|\Sigma|} \mathbb{1} \{s' = s \text{ or } B(v_{\sim j}[s], v_{\sim j}[s'])^c\} \tilde{q}(v_{\sim j}[s']) \\ &= \sum_{s' \in \Sigma \setminus \{s\}} \frac{1}{|\Sigma| - 1} \left(\frac{1}{|\Sigma|} \tilde{q}(v_{\sim j}[s]) + \frac{|\Sigma| - 1}{|\Sigma|} \mathbb{1} \{B(v_{\sim j}[s], v_{\sim j}[s'])^c\} \tilde{q}(v_{\sim j}[s']) \right). \end{aligned}$$

By Jensen's inequality, this implies

$$\log q(v_{\sim j}[s]) \geq \sum_{s' \in \Sigma \setminus \{s\}} \frac{1}{|\Sigma| - 1} \log \left(\frac{1}{|\Sigma|} \tilde{q}(v_{\sim j}[s]) + \frac{|\Sigma| - 1}{|\Sigma|} \mathbb{1} \{B(v_{\sim j}[s], v_{\sim j}[s'])^c\} \tilde{q}(v_{\sim j}[s']) \right).$$

Taking an expectation over $s \sim \text{unif}(\Sigma)$, we have

$$\begin{aligned} \mathbb{E}_{s \sim \text{unif}(\Sigma)} \log q(v_{\sim j}[s]) &\geq \sum_{\substack{s, s' \in \Sigma \\ s \neq s'}} \frac{1}{|\Sigma|(|\Sigma| - 1)} \left[\log \left(\frac{1}{|\Sigma|} \tilde{q}(v_{\sim j}[s]) + \frac{|\Sigma| - 1}{|\Sigma|} \mathbb{1} \{B(v_{\sim j}[s], v_{\sim j}[s'])^c\} \tilde{q}(v_{\sim j}[s']) \right) \right] \\ &= \sum_{\substack{s, s' \in \Sigma \\ s \neq s'}} \frac{1}{2|\Sigma|(|\Sigma| - 1)} \xi(v_{\sim j}, s, s'), \end{aligned} \tag{6.1}$$

where for $s \neq s'$,

$$\begin{aligned} \xi(v_{\sim j}, s, s') &= \log \left(\frac{1}{|\Sigma|} \tilde{q}(v_{\sim j}[s]) + \frac{|\Sigma| - 1}{|\Sigma|} \mathbb{1} \{B(v_{\sim j}[s], v_{\sim j}[s'])^c\} \tilde{q}(v_{\sim j}[s']) \right) \\ &\quad + \log \left(\frac{1}{|\Sigma|} \tilde{q}(v_{\sim j}[s']) + \frac{|\Sigma| - 1}{|\Sigma|} \mathbb{1} \{B(v_{\sim j}[s], v_{\sim j}[s'])^c\} \tilde{q}(v_{\sim j}[s]) \right). \end{aligned}$$

If $B(v_{\sim j}[s], v_{\sim j}[s'])$ holds, then $\xi(v_{\sim j}, s, s') = \log \tilde{q}(v_{\sim j}[s]) + \log \tilde{q}(v_{\sim j}[s']) - 2 \log |\Sigma|$. Otherwise, by Jensen's inequality we have

$$\log \left(\frac{1}{|\Sigma|} \tilde{q}(v_{\sim j}[s]) + \frac{|\Sigma| - 1}{|\Sigma|} \mathbb{1} \{B(v_{\sim j}[s], v_{\sim j}[s'])^c\} \tilde{q}(v_{\sim j}[s']) \right) \geq \frac{1}{|\Sigma|} \log \tilde{q}(v_{\sim j}[s]) + \frac{|\Sigma| - 1}{|\Sigma|} \log \tilde{q}(v_{\sim j}[s'])$$

and similarly for the other term of $\xi(v_{\sim j}, s, s')$. In this case, $\xi(v_{\sim j}, s, s') \geq \log \tilde{q}(v_{\sim j}[s]) + \log \tilde{q}(v_{\sim j}[s'])$. So, in all cases

$$\xi(v_{\sim j}, s, s') \geq \log \tilde{q}(v_{\sim j}[s]) + \log \tilde{q}(v_{\sim j}[s']) - 2 \mathbb{1} \{B(v_{\sim j}[s], v_{\sim j}[s'])\} \log |\Sigma|.$$

Substituting into (6.1), we have

$$\mathbb{E}_{s \sim \text{unif}(\Sigma)} \log q(v_{\sim j}[s]) \geq \mathbb{E}_{s \sim \text{unif}(\Sigma)} \log \tilde{q}(v_{\sim j}[s]) - \log |\Sigma| \cdot \sum_{\substack{s, s' \in \Sigma \\ s \neq s'}} \frac{\mathbb{1} \{B(v_{\sim j}[s], v_{\sim j}[s'])\}}{|\Sigma|(|\Sigma| - 1)}.$$

Taking an expectation over $v_{\sim j}$ yields

$$\mathbb{E}_{v \sim \text{unif}(G)} \log q(v) \geq \mathbb{E}_{v \sim \text{unif}(G)} \log \tilde{q}(v) - \log |\Sigma| \cdot \lambda_j.$$

By induction, we have

$$\mathbb{E}_{v \sim \text{unif}(G)} \log \tilde{q}(v) \geq -\log |\Sigma| \cdot \sum_{t=2}^T \lambda_{\sigma(t)},$$

and the result follows. $\square$

6.3 Completing the Proof of Stability

Proof of Proposition 6.2. Our interpolation scheme can be modeled as the random walk in Section 6.2, with $\Sigma = \mathcal{L}$, $J = km$, $T = k^2m$, $\sigma(t)$ defined in Definition 4.2, and where the bad edges are the c -bad edges. This correspondence is consistent because the steps in the interpolation path where a formula transitions to itself are never c -bad.

Since σ maps to every value in $[km]$ k times and $|\mathcal{L}| = 2n$, Proposition 6.3 and Lemma 6.5 imply that the probability of never traversing a c -bad edge is at least $(2n)^{-4Dk/c}$. $\square$

Proof of Proposition 4.7(a). Set $c = \frac{\beta_+ - \beta_-}{\gamma k}$. Let $S_{\text{no-bad}}$ be the event that for all $1 \leq t \leq T$, $(\Phi^{(t-1)}, \Phi^{(t)})$ is not c -bad with respect to f . By Proposition 6.2, $\mathbb{P}(S_{\text{no-bad}}) \geq (2n)^{-4Dk^2\gamma/(\beta_+ - \beta_-)}$.

By a union bound, $\mathbb{P}(S_{\text{valid}}) \geq 1 - (T+1)\delta$. Thus, $\mathbb{P}(S_{\text{valid}} \cap S_{\text{no-bad}}) \geq (2n)^{-4Dk^2\gamma/(\beta_+ - \beta_-)} - (T+1)\delta$. We claim that on $S_{\text{valid}} \cap S_{\text{no-bad}}$, the event S_{consec} also occurs.

Suppose for sake of contradiction that $S_{\text{valid}} \cap S_{\text{no-bad}}$ holds and for some $1 \leq t \leq T$, we have that $\Delta(x^{(t-1)}, x^{(t)}) > \frac{\beta_+ - \beta_-}{2k}$. Because $(\Phi^{(t-1)}, \Phi^{(t)})$ is not c -bad, we have

$$\left\| f(\Phi^{(t-1)}) - f(\Phi^{(t)}) \right\|_2^2 \leq c \mathbb{E}_{\Phi \sim \Phi_k(n, m)} \|f(\Phi)\|_2^2 \leq c\gamma n = \frac{\beta_+ - \beta_-}{k} n.$$

Let $I = \{i \in [n] : x_i^{(t-1)} \neq x_i^{(t)}\}$, so $|I| > \frac{\beta_+ - \beta_-}{2k} n$. Let $B^{(t-1)}$ and $B^{(t)}$ be the sets of indices B where the executions of $\mathcal{A}(x^{(t-1)})$ and $\mathcal{A}(x^{(t)})$ used the output of the assistance subroutine $\mathcal{B}$ (recall Definition 2.4). Because S_{valid} holds, the executions of $\mathcal{A}(x^{(t-1)})$ and $\mathcal{A}(x^{(t)})$ succeeded, and thus $|B^{(t-1)}|, |B^{(t)}| \leq \eta n = \frac{\beta_+ - \beta_-}{8k} n$.

Let $J = I \setminus (B^{(t-1)} \cup B^{(t)})$, so $|J| > \frac{\beta_+ - \beta_-}{4k} n$. For all $i \in J$, one of $f_i(\Phi^{(t-1)})$ and $f_i(\Phi^{(t)})$ is at least 1 and the other is at most -1 , and so $|f_i(\Phi^{(t-1)}) - f_i(\Phi^{(t)})| \geq 2$. So

$$\left\| f(\Phi^{(t-1)}) - f(\Phi^{(t)}) \right\|_2^2 \geq \sum_{i \in J} |f_i(\Phi^{(t-1)}) - f_i(\Phi^{(t)})|^2 > \frac{\beta_+ - \beta_-}{k} n.$$

This is a contradiction. Therefore $S_{\text{consec}} \supseteq S_{\text{valid}} \cap S_{\text{no-bad}}$, and so

$$\mathbb{P}(S_{\text{valid}} \cap S_{\text{consec}}) \geq \mathbb{P}(S_{\text{valid}} \cap S_{\text{no-bad}}) \geq (2n)^{-4Dk^2\gamma/(\beta_+ - \beta_-)} - (T+1)\delta.$$

$\square$

7 Proof of Achievability

In this section we will prove Theorem 2.10, that η -assisted low degree polynomial algorithms can solve random k -SAT at clause density $(1 - \varepsilon)2^k \log k/k$. We will achieve this by simulating (the first phase of) the algorithm **Fix** from [25] by a local algorithm on the factor graph, which we then simulate by a low degree polynomial.

This section is structured as follows. In Section 7.1, we define the k -SAT factor graph, formalize the notion of local algorithms, and state useful properties of these objects. The local algorithms we consider extend the *factors of i.i.d.* model considered in [51, 45, 69]. In Section 7.2, we introduce **Fix** and prove that it can be simulated by a local algorithm. In Section 7.3, we argue that any local algorithm can be simulated by a low degree polynomial, proving Theorem 2.10. Finally, in Section 7.4, we prove a technical proposition whose proof we deferred from Section 7.1.

7.1 Local Algorithms on the Factor Graph

We begin by introducing formalism for local algorithms on graphs. A (possibly infinite) graph $G = (V, E)$ is *locally finite* if every vertex $v \in V$ has a finite number of neighbors. Throughout this section, we will consider rooted decorated graphs, defined as follows.

Definition 7.1 (Rooted graph). A rooted graph is a pair (G, v) , where $G = (V, E)$ is a locally finite graph and $v \in V$ is a vertex of G .

As the definition suggests, we think of v as the root of G .

Definition 7.2 (Decorated graph). A decorated graph is a pair (G, σ) , where $G = (V, E)$ is locally finite and the *decoration map* $\sigma : E \rightarrow \{\mathbf{T}, \mathbf{F}\}$ assigns a boolean label to each edge of G .

Definition 7.3 (Rooted decorated graph). A rooted decorated graph is a triple (G, v, σ) , where (G, v) is a rooted graph and (G, σ) is a decorated graph. Let Λ denote the set of rooted decorated graphs.

We say $(G, v, \sigma), (G', v', \sigma') \in \Lambda$ are isomorphic if there is a root and decoration preserving isomorphism between them. We are now ready to define the two central graphs of our argument.

Definition 7.4 (k -SAT factor graph). Let $\mathcal{G}(n, m, k)$ denote the law of the factor graph of $\Phi_k(n, m)$, defined as follows. A sample $(G, \sigma) \sim \mathcal{G}(n, m, k)$ is a decorated bipartite graph with n left-vertices, representing variables, and m right-vertices, representing clauses. The left and right vertex sets are denoted $\text{Va}(G)$ and $\text{Cl}(G)$, respectively, and their union is denoted $V(G)$. There are k edges emanating from each $c \in \text{Cl}(G)$ to i.i.d. uniformly random (possibly repeated) vertices in $\text{Cl}(G)$. The edge set of G is denoted $E(G)$. The decoration map $\sigma : E(G) \rightarrow \{\mathbf{T}, \mathbf{F}\}$ decorates the edges i.i.d. from $\text{unif}(\{\mathbf{T}, \mathbf{F}\})$.

Each factor graph $(G, \sigma) \sim \mathcal{G}(n, m, k)$ encodes a k -SAT instance $\Phi \in \Omega_k(n, m)$: the k left-vertices adjacent to each right-vertex encode the k variables appearing in that clause, and the edge decorations σ encode the polarities of the literals.

Definition 7.5 (Decorated Alternating Poisson-Constant Galton-Watson Tree). Let $\text{DAGW}(d_1, d_2)$, with parameters $d_1 > 0, d_2 \in \mathbb{N}$, be the law of the following rooted decorated (possibly infinite) tree (T, o, σ) . First, we generate (T, o) by the following procedure.

- Start with a root vertex o in layer 0.
- For even ℓ , each vertex in layer ℓ independently spawns $\text{Pois}(d_1)$ children in layer $\ell + 1$. For odd ℓ , each vertex in layer ℓ spawns d_2 children in layer $\ell + 1$. Each non-root vertex is connected to its parent by an edge.

Let $V(T)$ and $E(T)$ denote the sets of vertices and edges of T , and let $\text{Va}(T)$ and $\text{Cl}(T)$ denote the sets of even and odd-depth vertices of T . Finally, $\sigma : E(T) \rightarrow \{\mathbf{T}, \mathbf{F}\}$ decorates the edges i.i.d. from $\text{unif}(\{\mathbf{T}, \mathbf{F}\})$.

The significance of this tree is that as $n \rightarrow \infty$ for fixed α, k , local neighborhoods of a fixed left-vertex of $\mathcal{G}(n, \alpha n, k)$ converge to local neighborhoods of the root of $\text{DAGW}(\alpha k, k - 1)$, in a sense formalized by Lemma 7.12 below. This is analogous to the fact that local neighborhoods of a fixed vertex of the sparse Erdős-Rényi graph $G(n, d/n)$ converge to local neighborhoods of the root of the Poisson Galton-Watson tree $\text{PGW}(d)$.

We will now build toward a definition of local algorithm. Informally, a local algorithm outputs a boolean value for each variable $v \in \text{Va}(G)$ of the k -SAT factor graph by looking at only a local neighborhood of v in the factor graph. We first formalize the notions of local neighborhood and local function.

Definition 7.6 (r -neighborhood). For a locally finite graph G , a vertex v of G , and a positive integer r , the r -neighborhood of v , denoted $N_r(G, v)$, is the rooted graph with root v containing all vertices of G reachable from v by a path of length at most r and all edges on these paths. We will denote the number of edges in $N_r(G, v)$ by $|N_r(G, v)|$.

Given a decoration map $\sigma : E(G) \rightarrow \{\mathbf{T}, \mathbf{F}\}$, we further define $N_r(G, v, \sigma)$ as the rooted decorated graph consisting of $N_r(G, v)$ with the decorations provided by σ on the edges therein.

Definition 7.7 (r -local function). A function g with domain Λ is r -local if $g(G, v, \sigma)$ depends only on the isomorphism class of $N_r(G, v, \sigma)$.

In other words, an r -local function decides its output by looking only at an r -neighborhood of the root, in which the root is distinguished but the remaining vertices are not.

We will allow our local algorithm to generate i.i.d. labels attached to each vertex and edge to assist its decision. So, let us formulate labeled variants of these definitions.

Definition 7.8. Let $G = (V, E)$ be a locally finite graph, and let $\varphi : V \cup E \rightarrow [0, 1]$. We will refer to φ as the *label map*.

- A *rooted decorated labeled graph* is a 4-tuple (G, v, σ, φ) where $(G, v, \sigma) \in \tilde{\Lambda}$. Let $\tilde{\Lambda}$ denote the set of rooted decorated labeled graphs. Two graphs $(G, v, \sigma, \varphi), (G', v', \sigma', \varphi') \in \tilde{\Lambda}$ are isomorphic if there is a root, decoration, and label preserving isomorphism between them.
- The *labeled r -neighborhood* $N_r(G, v, \sigma, \varphi)$ is the rooted decorated labeled graph consisting of $N_r(G, v)$ with the decorations and labels provided by σ, φ on the vertices and edges therein.
- A function h with domain $\tilde{\Lambda}$ is r -local if $h(G, v, \sigma, \varphi)$ depends only on the isomorphism class of $N_r(G, v, \sigma, \varphi)$.

We are now ready to define a local algorithm on the random k -SAT factor graph.

Definition 7.9 (Local Algorithm). Let $(G, \sigma) \sim \mathcal{G}(n, m, k)$, and let $h : \tilde{\Lambda} \rightarrow \{\mathbf{T}, \mathbf{F}\}$. The local algorithm based on h , which we denote $\mathcal{A}_h$, runs as follows on input (G, σ) .

- Generate labels $\varphi : V(G) \cup E(G) \rightarrow [0, 1]$, where each label is generated i.i.d. from $\text{unif}([0, 1])$.
- Output $x \in \{\mathbf{T}, \mathbf{F}\}^n$, where for each $v \in \text{Va}(G)$, $x_v = h(G, v, \sigma, \varphi)$.

We will abuse notation and identify (G, σ) with its corresponding k -SAT instance $\Phi \in \Omega_k(n, m)$. Thus we may treat Φ as an input to $\mathcal{A}_h$ as well.

Before we continue, we state four lemmas pertaining to local properties of graphs. Lemmas 7.10 and 7.11 control the growth rate of local neighborhoods of $\text{DAGW}(d_1, d_2)$ and $\mathcal{G}(n, m, k)$. Lemma 7.12 makes precise the sense in which local neighborhoods of left-vertices of $\mathcal{G}(n, \alpha n, k)$ converge to local neighborhoods of the root of $\text{DAGW}(\alpha k, k - 1)$. Lemma 7.13 gives that the sum of a local function concentrates. These lemmas are analogous to [12, Lemma 11.1, Lemma 11.2, Lemma 12.4 and Proposition 12.3], which give the analogous claims with $\mathcal{G}(n, m, k)$ and $\text{DAGW}(mk/n, k - 1)$ replaced by $G(n, d/n)$ and $\text{PGW}(d)$ (and for Lemmas 7.12 and 7.13, without the labels φ, φ' ; including these labels does not affect the results).

Lemma 7.10. Let $(T, o, \sigma) \sim \text{DAGW}(d_1, d_2)$ with $d_1, d_2 \geq 2$. There are two universal constants $c_0, c_1 > 0$ such that for every positive λ , we have

$$\mathbb{P}[|N_{2r}(T, o)| \leq \lambda(d_1 d_2)^r \text{ for all positive integers } r] \geq 1 - c_1 e^{-c_0 \lambda}.$$

Lemma 7.11. Let $(G, \sigma) \sim \mathcal{G}(n, m, k)$ with $\alpha = m/n$ and $\alpha k, k - 1 \geq 2$. Let $v \in \text{Va}(G)$ be a fixed vertex. There are two universal constants $c_0, c_1 > 0$ such that for every positive λ , we have

$$\mathbb{P}[|N_{2r}(G, v)| \leq \lambda(\alpha k(k - 1))^r \text{ for all positive integers } r] \geq 1 - c_1 e^{-c_0 \lambda}.$$

Lemma 7.12. Let $(G, \sigma) \sim \mathcal{G}(n, \alpha n, k)$ and let $(T, o, \sigma') \sim \text{DAGW}(\alpha k, k - 1)$. Let $\varphi : V(G) \cup E(G) \rightarrow [0, 1]$ and $\varphi' : V(T) \cup E(T) \rightarrow [0, 1]$ be labelings, where each label is drawn i.i.d. from $\text{unif}([0, 1])$. Let $g : \tilde{\Lambda} \rightarrow [-1, 1]$ be a $2r$ -local function. There is a universal constant $c > 0$ such that for all sufficiently large n (depending on α, k, r) and any fixed $v \in \text{Va}(G)$,

$$|\mathbb{E}[g(G, v, \sigma, \varphi)] - \mathbb{E}[g(T, o, \sigma', \varphi')]| \leq cn^{-0.49}.$$

Lemma 7.13. Let $(G, \sigma) \sim \mathcal{G}(n, \alpha n, k)$ with $\alpha k, k - 1 \geq 2$. Let $\varphi : V(G) \cup E(G) \rightarrow [0, 1]$ be a labeling, where each label is drawn i.i.d. from $\text{unif}([0, 1])$. Let $g : \tilde{\Lambda} \rightarrow [-1, 1]$ be a $2r$ -local function. There exists a universal constant $c > 0$ such that for all $p \geq 2$,

$$\mathbb{E} \left[\left| \sum_{v \in \text{Va}(G)} g(G, v, \sigma, \varphi) - \mathbb{E} \sum_{v \in \text{Va}(G)} g(G, v, \sigma, \varphi) \right|^p \right] \leq \left(cn^{1/2} p^{3/2} (\alpha k(k - 1))^r \right)^p.$$

The proofs of these lemmas can be easily adapted from the corresponding proofs of [12]. For the sake of brevity we only sketch the main ideas. Lemma 7.10 is proved by repeated conditioning, where we use a Chernoff bound to control the number of vertices at each odd depth conditioned on the number of vertices at the preceding even depth. Lemma 7.11 is proved by stochastic domination by $\text{DAGW}(\alpha k, k-1)$, followed by Lemma 7.10. Lemma 7.12 is proved by showing that $d_{\text{TV}}(N_{2r}(G, v, \sigma, \varphi), N_{2r}(T, o, \sigma', \varphi')) \leq cn^{-0.98}$, where d_{TV} denotes total variation distance, and then applying Cauchy-Schwarz. This total variation bound is in turn proved by a coupling of breadth-first search explorations of $N_{2r}(G, v, \sigma, \varphi)$ and $N_{2r}(T, o, \sigma', \varphi')$. where the key point is the estimate $d_{\text{TV}}(\text{Bin}(n, \lambda/n), \text{Pois}(\lambda)) \leq \lambda/n$. Lemma 7.13 is proved by [13, Theorem 15.5], where $Z = \sum_{v \in \text{Va}(G)} g(G, v, \sigma, \varphi)$ and the Z'_i are Z where we resample the k edges (and their decorations and labels) emanating from one clause $c \in \text{Cl}(G)$.

From Lemma 7.13, we can derive the following tail bound for sums of local functions.

Corollary 7.14. *Let $(G, \sigma) \sim \mathcal{G}(n, \alpha n, k)$ with $\alpha k, k-1 \geq 1$. Let $\varphi : V(G) \cup E(G) \rightarrow [0, 1]$ be a labeling, where each label is drawn i.i.d. from $\text{unif}([0, 1])$. Let $g : \tilde{\Lambda} \rightarrow [-1, 1]$ be a $2r$ -local function. There exists a universal constant $c > 0$ such that for all $t \geq (2e)^{3/2} cn^{1/2} (\alpha k(k-1))^r$,*

$$\mathbb{P} \left[\left| \sum_{v \in \text{Va}(G)} g(G, v, \sigma, \varphi) - \mathbb{E} \sum_{v \in \text{Va}(G)} g(G, v, \sigma, \varphi) \right| \geq t \right] \leq \exp \left(-\frac{3t^{2/3}}{2ec^{2/3}n^{1/3}(\alpha k(k-1))^{2r/3}} \right)$$

Proof. Let c be the constant from Lemma 7.13. Set $p = e^{-1} (cn^{1/2}(\alpha k(k-1))^r/t)^{-2/3} \geq 2$, so

$$\begin{aligned} & \mathbb{P} \left[\left| \sum_{v \in \text{Va}(G)} g(G, v, \sigma, \varphi) - \mathbb{E} \sum_{v \in \text{Va}(G)} g(G, v, \sigma, \varphi) \right| \geq t \right] \\ & \leq t^{-p} \mathbb{E} \left[\left| \sum_{v \in \text{Va}(G)} g(G, v, \sigma, \varphi) - \mathbb{E} \sum_{v \in \text{Va}(G)} g(G, v, \sigma, \varphi) \right|^p \right] \\ & \leq \exp \left(-\frac{3t^{2/3}}{2ec^{2/3}n^{1/3}(\alpha k(k-1))^{2r/3}} \right). \end{aligned}$$

□

Remark 7.15. Because we can choose $g : \tilde{\Lambda} \rightarrow [-1, 1]$ in Lemmas 7.12 and 7.13 and Corollary 7.14 that ignore the labels φ, φ' , these results also hold without the labels φ, φ' and with $g : \Lambda \rightarrow [-1, 1]$. We will sometimes apply these results in this form.

The following technical proposition will be useful to our analysis because it bounds the size of the largest (odd-depth) $2r$ -neighborhood in $N_{2t}(T, o)$ by a quantity that scales sublinearly in the depth t .

Proposition 7.16. *Let $d_1, d_2 \geq 2$ and $r \in \mathbb{N}$. For any $\eta \in (0, 1)$, there exist constants $C, t^* > 0$, depending on d_1, d_2, r, η , such that for all integers $t \geq t^*$ the following holds. If $(T, o, \sigma) \sim \text{DAGW}(d_1, d_2)$, then with probability at least $1 - \eta$, for all odd-depth vertices $c \in \text{Cl}(T) \cap N_{2t}(T, o)$, we have*

$$|N_{2r}(T, c)| \leq \frac{Ct}{\log^{(r)} t},$$

where $\log^{(r)}$ denotes the r th iterate of $\log$.

We defer the proof of this proposition to Section 7.4.

7.2 Simulating Fix with a Local Algorithm

At a high level, the algorithm **Fix** from [25] runs in three phases. In the first phase, the algorithm produces an almost-satisfying assignment. In the second phase, it modifies this assignment by changing a small number of variables to “don’t know.” This is done in a way such that the remaining problem of assigning truth values

to the “don’t know” variables is equivalent to a very subcritical 3-SAT instance. The third phase solves the remaining problem with a maxflow algorithm.

We will only simulate the first phase of **Fix**, which we will denote **Fix1**, by a local algorithm. This is enough to simulate **Fix** with slightly larger error because the second phase of **Fix** sets at most a k^{-12} fraction of bits to “don’t know” with high probability. Thus, a local algorithm that simulates **Fix1** with error η' simulates **Fix** with error $\eta = k^{-12} + \eta'$. Let us first formally record the guarantees on **Fix1** proved in [25].

Theorem 7.17 (Implicit in [25, Section 3]). *Let $\varepsilon > 0$. Let $\alpha = (1 - \varepsilon)2^k \log k/k$ and $m = \lfloor \alpha n \rfloor$. Let $\Phi \sim \Phi_k(n, m)$ and $x = \text{Fix1}(\Phi)$, where **Fix1** is defined below in Algorithm 7.18. With probability $1 - o(1)$, there exists a satisfying assignment y of Φ such that $\Delta(x, y) \leq k^{-12}$.*

Let us now define **Fix1**. To simplify notation, we name the variables $1, 2, \dots, n$. This phase starts from the all-true assignment and selects some variables $Z \subseteq [n]$ to set false so that most clauses are satisfied. To do this, it scans through the clauses of the input formula Φ . When it encounters an all-negative clause that does not contain any variable from Z , it tries to add a variable from this clause to Z in a way that does not create any more unsatisfied clauses. Formalizing this idea, we say a variable $v \in [n] \setminus Z$ is Z -safe if, when we set all variables in $[n] \setminus Z$ to true and all variables in Z to false, v is not the sole true literal in any clause.

Algorithm 7.18 (**Fix**, Phase 1). [25] **Fix1** takes as input $\Phi \in \Omega_k(n, m)$, and runs as follows.

- Set $Z = \emptyset$.
- Relabel the clauses $\Phi_1, \dots, \Phi_m$ in a uniformly random order. Also, for each Φ_i , relabel the literals $\Phi_{i,1}, \dots, \Phi_{i,k}$ in a uniformly random order.
- For each $i = 1, \dots, m$,
 - If Φ_i is all-negative and contains no variable from Z :
 - * If there is $1 \leq j < \lceil k/2 \rceil$ such that the underlying variable of $\Phi_{i,j}$ is Z -safe, pick the smallest such j and add the underlying variable of $\Phi_{i,j}$ to Z .
 - * Otherwise, add the underlying variable of $\Phi_{i,\lceil k/2 \rceil}$ to Z .
- Output $x \in \{\text{T}, \text{F}\}^n$ where $x_i = \text{F}$ if $i \in Z$ and otherwise $x_i = \text{T}$.

The presentation of **Fix1** in [25] does not rerandomize the clause and literal orders, but of course this makes no difference. For technical reasons having to do with the analysis in [25], **Fix1** only considers flipping variables $\Phi_{i,j}$ where $j \leq \lceil k/2 \rceil$. The main result of this subsection is the following proposition, which shows that **Fix1** can be simulated by a local algorithm.

Proposition 7.19 (Simulating **Fix1** with local algorithm). *Let α, k, n, m be as in Theorem 7.17, with $\alpha k, k - 1 \geq 2$. For any $\eta > 0$, there exists a positive integer R (depending on η) and an R -local function $h : \tilde{\Lambda} \rightarrow \{\text{T}, \text{F}\}$, such that, for some coupling of the internal randomnesses of **Fix1** and $\mathcal{A}_h$, and for $\Phi \sim \Phi_k(n, m)$, we have*

$$\mathbb{P}[\Delta(\text{Fix1}(\Phi), \mathcal{A}_h(\Phi)) \geq \eta] \leq \exp(-\Omega(n^{1/3})).$$

We first sketch informally why such simulation should be possible. While **Fix1** is not a local algorithm, it is “sequentially local,” in the following sense.³ Like a local algorithm, **Fix1** makes decisions that depend only on information available in a local neighborhood. Namely, the logic inside the for loop depends only on a 3-neighborhood of the clause Φ_i on the factor graph. But unlike a true local algorithm, which makes all of its decisions in parallel, **Fix1** makes its decisions in series, and each decision may leave information on the vertices it accessed which future decisions can see. In a sense that will be made precise later, sequentiality cannot induce long dependence chains. So, we can simulate **Fix** by simulating it on a (larger) local neighborhood of each variable.

To prove Proposition 7.19, we will define the class of *sequentially local recording algorithms* and show the stronger claim that any such algorithm in this class can be simulated by a local algorithm.

³**Fix1** is “sequentially local” in a slightly different sense than the sequential local algorithms considered in [46], because it sequentially makes local decisions on neighborhoods of each *clause*, while the sequentially local algorithms in [46] sequentially make local decisions on neighborhoods of each *variable*. Nonetheless, by a nearly identical argument we can show that the latter algorithms can also be simulated by local algorithms. The proof idea remains the same: sequentiality cannot induce long dependence chains.

Definition 7.20 (Sequentially local recording algorithm). Let $(G, \sigma) \sim \mathcal{G}(n, m, k)$, and let r be a positive integer. A sequentially r -local recording algorithm receives input (G, σ) and runs as follows.

- Initialize the memory map $\mu : \text{Va}(G) \rightarrow \mathbb{N}$ as the all-0 map.
- Generate labels $\psi : E(G) \rightarrow [0, 1]$ i.i.d. from $\text{unif}([0, 1])$.
- For clauses $c \in \text{Cl}(G)$, in a uniformly random order from the $m!$ possible permutations of $\text{Cl}(G)$:
 - Run a procedure $h_1(G, \sigma, \mu, \psi)$, which depends only on $N_r(G, c)$ and the memory μ and labels ψ associated with variables (left-vertices) and edges in this neighborhood. This procedure may overwrite the values of $\mu(v)$ for any $v \in \text{Va}(G) \cap N_r(G, c)$.
- Output $x \in \{\text{T}, \text{F}\}^n$, where for each variable $v \in \text{Va}(G)$, we set $x_v = h_2(\sigma(v))$ for a deterministic procedure σ .

As in Definition 7.9, we will abuse notation and treat $\Phi \in \Omega_k(n, m)$, which we identify with its factor graph (G, σ) , as an input to a sequentially local recording algorithm.

Informally, these algorithms have access to a memory map $\mu : V(G) \rightarrow \mathbb{N}$, which we think of as an unlimited notepad on each variable, and internally generated randomness ψ for each edge. It processes clauses $c \in \text{Cl}(G)$ in a uniformly random order; in each step, the subroutine h_1 accesses the r -neighborhood of c and can overwrite the data written on any variable in that neighborhood. In the end, each variable decides to be true or false depending on the final value written on its notepad.

Let us see that **Fix1** is in this class.

Fact 7.21. *There exists a sequentially 3-local recording algorithm $\mathcal{B}$ such that the internal randomnesses of $\mathcal{B}$ and **Fix1** can be coupled so that for all $\Phi \in \Omega_k(n, m)$, $\mathcal{B}(\Phi) = \text{Fix1}(\Phi)$ almost surely.*

Proof. We couple the clause orderings of **Fix1** and $\mathcal{B}$ so that their for loops run over the clauses in the same order. For each clause $c \in \text{Cl}(G)$, we couple the rerandomization of the literal order within c in **Fix1** with the edge labels $\{\psi(e) : e \text{ incident to } c\}$ so that **Fix1** orders these literals in increasing order of the labels on their edges to c .

We will simulate Z with μ : over the coupled executions of **Fix1** and $\mathcal{B}$, we will keep the invariant that for each $v \in \text{Va}(G)$, $\mu(v) = 1$ if $v \in Z$, and $\mu(v) = 0$ otherwise. We can easily verify that the logic inside the for loop only depends on a 3-neighborhood of Φ_i . $\square$

We now define the local algorithm that will simulate an r -local recording algorithm and state our simulation result, Proposition 7.23. Intuitively, the R -local simulation $\mathcal{A}$ of $\mathcal{B}$ runs $\mathcal{B}$ restricted to $N_R(G, v)$, and because sequentiality cannot induce long-range dependencies this will often give the correct output.

Definition 7.22 (R -local simulation). Suppose $\mathcal{B}$ is a sequentially r -local recording algorithm. For a positive integer R , the R -local simulation of $\mathcal{B}$ is the R -local algorithm $\mathcal{A} = \mathcal{A}_h$, whose internal randomness is coupled with that of $\mathcal{B}$ as follows.

- The labels $\{\varphi(e) : e \in E(G)\}$ generated by $\mathcal{A}$ equal the labels $\{\psi(e) : e \in E(G)\}$ generated by $\mathcal{B}$.
- The clause ordering of $\mathcal{B}$ is coupled with the labels $\{\varphi(c) : c \in \text{Cl}(G)\}$ generated by $\mathcal{A}$, so that the for loop of $\mathcal{B}$ runs through the clauses in increasing order of $\varphi(c)$.
- The labels $\{\varphi(v) : v \in \text{Va}(G)\}$ are independent of the internal randomness of $\mathcal{B}$ (and unused by $\mathcal{A}$).

The R -local function $h : \tilde{\Lambda} \rightarrow [0, 1]$ associated with $\mathcal{A}$ is as follows. On input (G, v, σ, φ) , h treats $N_R(G, v, \sigma)$ as a factor graph, where the even-depth vertices represent variables and the odd-depth vertices represent clauses.⁴ h runs $\mathcal{B}$ on this factor graph, where the labels ψ are those of $\{\varphi(e) : e \in E(G) \cap N_R(G, v)\}$ and the for loop runs through the clauses in increasing order of $\varphi(c)$ for $c \in \text{Cl}(G) \cap N_R(G, v)$.

⁴Even and odd depth are well defined because G is bipartite.

Proposition 7.23 (Local algorithms simulate sequentially local recording algorithms). *Suppose $k\alpha, k-1 \geq 2$. Let $\mathcal{B}$ be a sequentially r -local recording algorithm. For any $\eta > 0$, there exists a positive integer R (depending on η) such that the R -local simulation $\mathcal{A}$ of $\mathcal{B}$ satisfies*

$$\mathbb{P}[\Delta(\mathcal{A}(\Phi), \mathcal{B}(\Phi)) \geq \eta] \leq \exp(-\Omega(n^{1/3})).$$

From these results, Proposition 7.19 is immediate.

Proof of Proposition 7.19. This follows from Fact 7.21 and Proposition 7.23. $\square$

It remains to prove Proposition 7.23. The main idea of the proof is that if, in the execution of $\mathcal{B}$, the output on v is influenced by a chain of t consecutive decisions, the following structure must exist in its local simulations.

Definition 7.24 ((r, t) -dependence chain). Let r, t be positive integers. Let (G, v) be a bipartite rooted graph with vertices $V(G)$, edges $E(G)$, and odd-depth vertices $\text{Cl}(G)$. Let $\varphi : V(G) \cup E(G) \rightarrow [0, 1]$ be a labeling. An (r, t) -dependence chain of (G, v, φ) is a sequence $(c_1, \dots, c_t) \in \text{Cl}(G)$ such that $c_1 \in N_r(v)$, $c_{i+1} \in N_{2r}(c_i)$ for $i = 1, \dots, t-1$, and the labels $\varphi(c_1), \dots, \varphi(c_t)$ form a decreasing sequence.

We will first bound the probability that this structure arises in $\text{DAGW}(d_1, d_2)$. We will then translate this to a bound on the number of times this structure arises in the k -SAT factor graph, via the machinery in Lemma 7.12 and Corollary 7.14.

Proposition 7.25. *Let $d_1, d_2 \geq 2$, $r \in \mathbb{N}$, and $\eta \in (0, 1)$. Let $(T, o, \sigma) \sim \text{DAGW}(d_1, d_2)$, and let $\varphi : V(T) \cup E(T) \rightarrow [0, 1]$ be labels generated i.i.d. from $\text{unif}([0, 1])$. There exists $t^* > 0$, depending on d_1, d_2, r, η , such that for all integers $t \geq t^*$, (T, o, φ) does not have an (r, t) -dependence chain with probability at least $1 - \eta$.*

Proof. By Lemma 7.10, we can set $\lambda > 0$ such that with probability $1 - \eta/3$, $|N_{2r}(T, o)| \leq \lambda(d_1 d_2)^r$. By applying Proposition 7.16 with $\eta/3$ in place of η and $R = rt$ in place of t , we get that (for sufficiently large t) with probability $1 - \eta/3$, for all $c \in \text{Cl}(T) \cap N_{2rt}(T, o)$, we have $|N_{2r}(T, c)| \leq \frac{Ct}{\log^{(r)} t}$, for another constant C . By a union bound, both events occur with probability $1 - 2\eta/3$.

On the intersection of these events, the number of sequences $(c_1, \dots, c_t) \in \text{Cl}(G)$ such that $c_1 \in N_r(v)$, $c_{i+1} \in N_{2r}(c_i)$ for $i = 1, \dots, t-1$, is at most

$$\lambda(d_1 d_2)^r \left(\frac{Ct}{\log^{(r)} t} \right)^{t-1}.$$

This is because there are at most $\lambda(d_1 d_2)^r$ choices for c_1 , and at most $\frac{Ct}{\log^{(r)} t}$ choices of c_{i+1} given c_i . In each of these sequences, the labels $\varphi(c_1), \dots, \varphi(c_t)$ form a decreasing sequence with probability $\frac{1}{t!}$, so the expected number of (r, t) dependence chains is at most

$$\lambda(d_1 d_2)^r \left(\frac{Ct}{\log^{(r)} t} \right)^{t-1} \cdot \frac{1}{t!}.$$

Because $t! \sim \sqrt{2\pi t} \left(\frac{t}{e}\right)^t$, for a large enough constant t this number is at most $\eta/3$. By Markov's inequality, for this t , the probability of there being an (r, t) -dependence chain is at most $\eta/3$.

By a final union bound, the result follows. $\square$

Finally, we can prove Proposition 7.23.

Proof of Proposition 7.23. We set t large enough that Proposition 7.25 holds with $k\alpha, k-1, r, \eta/3$ in place of d_1, d_2, r, η . Let $(T, o, \sigma) \sim \text{DAGW}(k\alpha, k-1)$ and $\varphi' : V(T) \cup E(T) \rightarrow [0, 1]$ be labels generated i.i.d. from $\text{unif}([0, 1])$. Thus with probability $1 - \eta/3$, (T, o, φ') does not have an (r, t) -dependence chain.

Let (G, σ) be the factor graph of Φ . We set $R = 2rt$. Let $\mathcal{A}$ be the R -local simulation of $\mathcal{B}$ and $\varphi : V(G) \cup E(G) \rightarrow [0, 1]$ be the internal randomness of $\mathcal{A}$. For $v \in \text{Va}(G)$, if (G, v, φ) does not have a (t, r) -dependence chain, then $\mathcal{A}$ and $\mathcal{B}$ output the same bit in the v position. Moreover, the event that

(G, v, φ) has a (t, r) -dependence chain is an R -local function of (G, v, σ, φ) . Thus, applying Corollary 7.14 with $t = \frac{\eta}{3}n$ and Lemma 7.12, we have that

$$\begin{aligned} \Delta(\mathcal{A}(\Phi), \mathcal{B}(\Phi)) &\leq \frac{1}{n} \sum_{v \in \text{Va}(G)} \mathbb{1} \{(G, v, \varphi) \text{ has a } (t, r)\text{-dependence chain}\} \\ &\leq \frac{\eta}{3} + \mathbb{E} \mathbb{1} \{(G, v, \varphi) \text{ has a } (t, r)\text{-dependence chain}\} \\ &\leq \frac{\eta}{3} + cn^{-0.49} + \mathbb{E} \mathbb{1} \{(T, o, \varphi') \text{ has a } (t, r)\text{-dependence chain}\} \\ &\leq \frac{2\eta}{3} + cn^{-0.49} < \eta, \end{aligned}$$

where the second inequality occurs with probability $1 - \exp(-\Omega(n^{1/3}))$ and the remaining steps are deterministic. $\square$

7.3 Simulating a Local Algorithm by a Low Degree Polynomial

So far, we have showed that **Fix1** can be simulated by a local algorithm. In this section, we will show that any local algorithm can be simulated by a low degree polynomial, thereby completing the proof of Theorem 2.10.

To formalize parsing the output of a polynomial as a boolean assignment, we introduce the symbol $?$ and define the function $\text{boolify} : \mathbb{R} \rightarrow \{\mathbf{T}, \mathbf{F}, ?\}$ by

$$\text{boolify}(x) = \begin{cases} \mathbf{T} & x = 1, \\ \mathbf{F} & x = -1, \\ ? & \text{otherwise.} \end{cases}$$

When applied to a vector, boolify is applied coordinate-wise. Note that this is a more demanding parsing scheme than the one in Definition 2.4. We will show in Proposition 7.26 that we can in fact simulate a local algorithm with a low degree polynomial, parsed by this stronger rule.

Let $N = m \cdot k \cdot 2n$, and recall that each $\Phi \in \Omega_k(n, m)$ can be identified as a vector in $\{0, 1\}^N$, as described in Section 2.

Proposition 7.26. *Let $\mathcal{A} = \mathcal{A}_h$ be an r -local algorithm, where $h : \tilde{\Lambda} \rightarrow \{\mathbf{T}, \mathbf{F}\}$ is an r -local function. For any $\eta > 0$, there exist constants $D, \gamma > 0$ (depending on r, η) and a (random) degree- D polynomial $f : \mathbb{R}^N \times \Omega \rightarrow \mathbb{R}^n$, whose randomness $\omega \sim \Omega$ is coupled to the internal randomness of $\mathcal{A}$, such that the following properties hold.*

- $\mathbb{P}[\Delta(\mathcal{A}_h(\Phi), \text{boolify}(f(\Phi))) \geq \eta] \leq \exp(-\Omega(n^{1/3}))$.
- $\mathbb{E}_{\omega, \Phi} \|f(\Phi)\|_2^2 \leq \gamma n$.

The proof of this proposition closely resembles the proof of [71, Theorem 1.4].

Proof. Let (G, σ) be the factor graph corresponding to $\Phi \sim \Phi_k(n, m)$; we identify the set $\text{Va}(G)$ with $[n]$. The internal randomness of f samples the same $\varphi : V(G) \cup E(G) \rightarrow [0, 1]$ as $\mathcal{A}$. Let s be a constant (depending on r, η) to be chosen later. We will construct f with the following property:

$$\text{For all } v \in \text{Va}(G), \text{ if } N_r(G, v) \text{ is a tree and } |N_r(G, v)| \leq s, \text{ then } f_v(\Phi, \varphi) = \text{boolify}^{-1}(h(G, v, \sigma, \varphi)). \quad (7.1)$$

We construct f recursively, as follows. Recall that the vectorization of Φ consists of indicators $\Phi_{i,j,\ell}$ of the events that the j th literal of the i th clause of Φ is the ℓ th literal of $\mathcal{L}$. We can naturally associate the triple (i, j, ℓ) with an edge of the factor graph, where ℓ determines the edge's endpoint on $\text{Va}(G)$ and i determines the endpoint of the edge in $\text{Cl}(G)$; denote this edge $e(i, j, \ell)$. For a set $S \subseteq [m] \times [k] \times [2n]$, define $e(S)$ as the (multi-)graph on vertex set $V(G)$ with edges $e(i, j, \ell)$ for each $(i, j, \ell) \in S$. Let $\mathcal{G}_{v,r,\ell}$ be the collection of sets $S \subseteq [m] \times [k] \times [2n]$, such that

- (a) $e(S)$ is a simple graph, and a tree in which every non-isolated vertex has a path to v of length at most r .
- (b) $|S| \leq s$.

Equivalently, $\mathcal{G}_{v,r,s}$ is the set of collections of (i, j, ℓ) corresponding to all possible r -neighborhoods of v of size at most s . We will set

$$f_v(\Phi, \varphi) = \sum_{S \in \mathcal{G}_{v,r,s}} \alpha(e(S), v, \sigma, \varphi) \prod_{(i,j,\ell) \in S} \Phi_{(i,j,\ell)}. \quad (7.2)$$

where the coefficients $\alpha(e(S), v, \sigma, \varphi)$ are chosen so the above property is satisfied. That is, $\alpha(e(S), v, \sigma, \varphi)$ is defined recursively by

$$\alpha(e(S), v, \sigma, \varphi) = \text{boolify}^{-1}(h(e(S), v, \sigma, \varphi)) - \sum_{S' \subsetneq S} \alpha(e(S'), v, \sigma, \varphi). \quad (7.3)$$

It is clear that this f satisfies (7.1). To show the first conclusion, we will show that with probability at least $1 - \exp(-\Omega(n^{1/3}))$, the number of $v \in \text{Va}(G)$ where $N_r(G, v)$ is a tree and $|N_r(G, v)| \leq s$ is at least $(1 - \eta)n$.

Let $\mathcal{T}$ be the set of rooted trees containing one representative from each isomorphism class of rooted trees of depth at most r . Let $\mathcal{T}_s$ be the subset of $\mathcal{T}$ containing the trees with at most s edges. For $T \in \mathcal{T}$, let n_T be the number of vertices $v \in \text{Va}(G)$ with r -neighborhood T , i.e.

$$n_T = \#(v \in \text{Va}(G) : N_r(G, v) \cong T).$$

Similarly, let p_T be the probability T occurs as the r -neighborhood of the root in $\text{DAGW}(k\alpha, k-1)$, i.e.

$$p_T = \mathbb{P}_{(U,o,\sigma) \sim \text{DAGW}(k\alpha, k-1)} [N_r(U, o) = T]$$

Choose s large enough that $\sum_{T \in \mathcal{T}_s} p_T \geq 1 - \eta/3$. Because the function $(G, v) \mapsto \mathbf{1}\{N_r(G, v) \cong T\}$ is r -local, Lemma 7.12 gives that for each $T \in \mathcal{T}_s$,

$$|\mathbb{E}[n_T] - p_T n| \leq cn^{0.51}.$$

Corollary 7.14, with $t = \frac{\eta n}{3|\mathcal{T}_s|}$, gives that with probability $1 - \exp(-\Omega(n^{1/3}))$,

$$|n_T - \mathbb{E}[n_T]| \leq \frac{\eta n}{3|\mathcal{T}_s|}.$$

Suppose this event occurs. Then,

$$\begin{aligned} \sum_{T \in \mathcal{T}_s} n_T &\geq \sum_{T \in \mathcal{T}_s} \left(\mathbb{E}[n_T] - \frac{\eta n}{3|\mathcal{T}_s|} \right) \geq \sum_{T \in \mathcal{T}_s} \left(p_T n - cn^{0.51} - \frac{\eta n}{3|\mathcal{T}_s|} \right) = n \sum_{T \in \mathcal{T}_s} p_T - \frac{\eta n}{3} - c|\mathcal{T}_s|n^{0.51} \\ &\geq \left(1 - \frac{2\eta}{3} \right) n - c|\mathcal{T}_s|n^{0.51} \geq (1 - \eta)n. \end{aligned}$$

This proves the first conclusion.

To prove the second conclusion, we will show that $\mathbb{E}_{\Phi, \varphi}[f_v(\Phi, \varphi)^2] = O(1)$ uniformly over $v \in \text{Va}(G)$. Define the random variable $M = |N_r(G, v)|$. In the expansion (7.2), the monomial indexed by $S \in \mathcal{G}_{v,r,s}$ is only nonzero if $e(S)$ is a subgraph of $N_r(G, v)$. So, the number of nonzero monomials is at most

$$k^s \sum_{i=0}^s \binom{M}{i} \leq k^s (M+1)^s.$$

Moreover, from (7.3) we see that each of the coefficients $\alpha(e(S), v, \sigma, \varphi)$ is upper bounded by a constant a dependent on r, s . So,

$$f_v(\Phi, \varphi)^2 \leq a^2 k^{2s} (M+1)^{2s}.$$

Lemma 7.11 gives that $\mathbb{P}[M \geq x] \leq C_0 \exp(-C_1 x)$ for some constants C_0, C_1 ; by integration by tails, we get $\mathbb{E}[(M+1)^{2s}] = O(1)$. So, $\mathbb{E}[f_v(\Phi, \varphi)^2] = O(1)$, as desired. $\square$

The proof of Theorem 2.10 follows readily from this proposition and the previous results.

Proof of Theorem 2.10. Set $\eta' = \frac{\eta - k^{-12}}{2}$. By Proposition 7.19, there exists a positive integer R and R -local function $h : \tilde{\Lambda} \rightarrow \{\mathbf{T}, \mathbf{F}\}$ such that for some coupling of the internal randomness of $\mathbf{Fix1}$ and $\mathcal{A}_h$,

$$\mathbb{P}[\Delta(\mathbf{Fix1}(\Phi), \mathcal{A}_h(\Phi)) \geq \eta'] \leq \exp(-\Omega(n^{1/3})).$$

Let f be the random degree- D polynomial given by Proposition 7.26. For some coupling of the internal randomness of $\mathcal{A}_h$ and f ,

$$\mathbb{P}[\Delta(\mathcal{A}_h(\Phi), \text{boolify}(f(\Phi))) \geq \eta'] \leq \exp(-\Omega(n^{1/3})).$$

Combined with Theorem 7.17, we have that except with probability $\delta = o(1) + \exp(-\Omega(n^{1/3})) = o(1)$, $\text{boolify}(f(\Phi))$ outputs an assignment within normalized Hamming distance $k^{-12} + 2\eta' = \eta$ of a satisfying assignment of Φ . Proposition 7.26 also gives that $\mathbb{E}_{\omega, \Phi} \|f(\Phi)\|_2^2 \leq \gamma n$.

Thus f is a random degree- D polynomial that (δ, γ, η) -solves $\Phi_k(n, m)$. Finally, by Lemma 4.1, we can make f deterministic in exchange for a constant factor in δ, γ . $\square$

7.4 Proof of Proposition 7.16

In this subsection, we will prove the technical Proposition 7.16. We introduce the following inverted version of $\text{DAGW}(d_1, d_2)$, which will be useful because the neighborhoods of odd-depth vertices of $\text{DAGW}(d_1, d_2)$ resemble $\text{IDAGW}(d_1, d_2)$, in a way formalized by Fact 7.28 below.

Definition 7.27 (Inverted Decorated Alternating Galton-Watson Tree). Let $\text{IDAGW}(d_1, d_2)$, with parameters $d_1 > 0, d_2 \in \mathbb{N}$ be the law of the following rooted decorated (possibly infinite) tree (T, o, σ) . The root o has $d_2 + 1$ children each connected to o by an edge, and σ decorates each of these edges independently from $\text{unif}(\{\mathbf{T}, \mathbf{F}\})$. The descendant subtrees of each of these children are i.i.d. copies of $\text{DAGW}(d_1, d_2)$.

Fact 7.28. Let $(T, o, \sigma) \sim \text{DAGW}(d_1, d_2)$ and $(T', o', \sigma') \sim \text{IDAGW}(d_1, d_2)$, and let v be an odd-depth vertex of T . Then, for any r , $|N_r(T, v)|$ is stochastically dominated by $|N_r(T', o')|$.

Proof. When we re-root (T, o) at v , we get a sample from $\text{IDAGW}(d_1, d_2)$, except that the original root o has one fewer child. $\square$

Next, we prove a variant of Lemma 7.10 for IDAGW and *finite* r , where we achieve a better error probability far in the tail.

Lemma 7.29. Let $d_1, d_2 \geq 2$ and $r \in \mathbb{N}$. Let $(T, o, \sigma) \sim \text{IDAGW}(d_1, d_2)$. Then, for all sufficiently large s (depending on d_1, d_2, r) we have that

$$\mathbb{P}\left[|N_{2r}(T, o)| > \frac{2s}{\log^{(r)} s} (d_1 d_2)^r\right] \leq e^{-s}.$$

Proof. For $1 \leq \ell \leq 2r$, let S_ℓ denote the number of vertices in $\text{IDAGW}(d_1, d_2)$ at depth ℓ . Then, $S_1 = d_2 + 1$. For odd $\ell \geq 1$, $S_{\ell+1}$ is the sum of S_ℓ i.i.d. copies of $\text{Pois}(d_1)$, and for even $\ell \geq 1$, $S_{\ell+1} = d_2 S_\ell$.

For $u = 1, 2, \dots, r$, define the events

$$E_u = \left\{S_{2u-1} \leq \frac{s}{\log^{(u-1)} s} (d_2 + 1)(d_1 d_2)^{u-1}\right\} \quad \text{and} \quad E'_u = \left\{S_{2u} \leq \frac{s}{\log^{(u)} s} d_1 (d_2 + 1)(d_1 d_2)^{u-1}\right\}.$$

On $\bigcap_{u=1}^r (E_u \cap E'_u)$, we have

$$|N_{2r}(T, o)| \leq \sum_{\ell=1}^{2r} S_\ell \leq \frac{s}{\log^{(r)} s} \cdot \frac{(d_1 + 1)(d_2 + 1)(d_1 d_2)^{r-1}}{1 - (d_1 d_2)^{-1}} \leq \frac{2s}{\log^{(r)} s} (d_1 d_2)^r.$$

So, it remains to show that $\mathbb{P}[\bigcap_{u=1}^r (E_u \cap E'_u)] \geq 1 - e^{-s}$. Note that E_1 holds by definition, while E'_u is equivalent to E_{u+1} . We will upper bound $\mathbb{P}((E'_u)^c | E_u)$. Let $N = \frac{s}{\log^{(u-1)} s} (d_2 + 1)(d_1 d_2)^{u-1}$. Conditioned on E_u , S_{2u} is stochastically dominated by $\sum_{i=1}^N x_i$, where the x_i are i.i.d. samples from $\text{Pois}(d_1)$. So,

$$\begin{aligned} \mathbb{P}[(E'_u)^c | E_u] &\leq \mathbb{P}\left[\sum_{i=1}^N x_i > \frac{\log^{(u-1)} s}{\log^{(u)} s} d_1 N\right] \leq \left(\inf_{s>0} \frac{\mathbb{E} \exp(sx_1)}{\exp\left(\frac{\log^{(u-1)} s}{\log^{(u)} s} d_1\right)}\right)^N = \exp\left(-Nd_1 \gamma\left(\frac{\log^{(u-1)} s}{\log^{(u)} s}\right)\right) \\ &= \exp\left(-\frac{s}{\log^{(u-1)} s} d_1 (d_2 + 1)(d_1 d_2)^{u-1} \gamma\left(\frac{\log^{(u-1)} s}{\log^{(u)} s}\right)\right). \end{aligned}$$

where $\gamma(x) = x \log x - x + 1$. We can take s large enough that for all $u = 1, \dots, r$, we have

$$\gamma\left(\frac{\log^{(u-1)} s}{\log^{(u)} s}\right) \geq \frac{1}{2} \log^{(u-1)} s,$$

while the bounds $d_1, d_2 \geq 2$ give $d_1(d_2 + 1)(d_1 d_2)^{u-1} \geq 2(u + 1)$. Thus,

$$\mathbb{P}[(E'_u)^c | E_u] \leq \exp(-(u + 1)s),$$

and so

$$\mathbb{P}\left[\bigcap_{u=1}^r (E_u \cap E'_u)\right] \geq 1 - \sum_{u=1}^r \mathbb{P}[(E'_u)^c | E_u] \geq 1 - \sum_{u=1}^{\infty} e^{-(u+1)s} \geq 1 - e^{-s}$$

for sufficiently large s . □

We can now complete the proof of Proposition 7.16.

Proof of Proposition 7.16. Set constant $\lambda > 0$ such that the conclusion of Lemma 7.10 holds with probability $1 - \eta/2$. On this event, $|N_{2t}(T, o)| \leq \lambda(d_1 d_2)^t$ for all t . Set s so that $\lambda(d_1 d_2)^t e^{-s} = \eta/2$; note that $s = \Theta(t)$. For all sufficiently large t , s is large enough that Lemma 7.29 applies. Then, by Fact 7.28, Lemma 7.29, and a union bound, $|N_{2r}(T, c)| \leq \frac{2s}{\log^{(r)} s} (d_1 d_2)^r$ for all $c \in \text{Cl}(T) \cap N_{2t}(T, o)$ with probability $1 - \eta/2$. Because $s = O(t)$, we have $\frac{2s}{\log^{(r)} s} (d_1 d_2)^r \leq \frac{Ct}{\log^{(r)} t}$ for some C . By a union bound, the conclusion occurs with probability $1 - \eta$. □

References

- [1] Dimitris Achlioptas. Random Satisfiability. *Handbook of Satisfiability*, 185:245-270, 2009.
- [2] Dimitris Achlioptas, Paul Beame, and Michael Molloy. Exponential bounds for DPLL below the satisfiability threshold. *Proceedings of 15th SODA*, 139-140, 2004.
- [3] Dimitris Achlioptas and Amin Coja-Oghlan. Algorithmic barriers from phase transitions. *Proceedings of 49th FOCS*, 793-802, 2008.
- [4] Dimitris Achlioptas and Gregory B. Sorkin. Optimal myopic algorithms for random 3-SAT. *Proceedings of 41st FOCS*, 590-600, 2000.
- [5] Miklós Ajtai. Σ_1^1 -formulae on finite structures. *Annals of Pure and Applied Logic*, 24(1):1-48, 1983.
- [6] Gérard B. Arous, Alexander S. Wein, and Ilias Zadik. Free energy wells and overlap gap property in sparse PCA. *Proceedings of 33rd COLT*, 479-482, 2020.
- [7] Afonso S. Bandeira, Dmitriy Kunisky, and Alexander S. Wein. Computational hardness of certifying bounds on constrained PCA problems. *Proceedings of 11th ITCS*, 2020.

- [8] Afonso S. Bandeira, Jess Banks, Dmitriy Kunisky, Cristopher Moore, and Alexander S. Wein. Spectral planting and the hardness of refuting cuts, colorability, and communities in random graphs. *Conference on Learning Theory (COLT)*, 2021.
- [9] Boaz Barak, Samuel B. Hopkins, Jonathan Kelner, Pravesh K. Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM Journal on Computing*, 48(2):687-735, 2019.
- [10] Mohsen Bayati, David Gamarnik, and Prasad Tetali. Combinatorial approach to the interpolation method and scaling limits in sparse random graphs. *Proceedings of 42nd STOC*, 105-114, 2010.
- [11] Mohsen Bayati and Andrea Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57(2):764-785, 2011.
- [12] Charles Bordenave, Simon Coste, and Raj Rao Nadakuditi. Detection thresholds in very sparse matrix completion. *arXiv:2005.06062*, preprint 2020.
- [13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities*. Oxford University Press, 2013.
- [14] Alfredo Braunstein, Marc Mézard, and Riccardo Zecchina. Survey propagation: an algorithm for satisfiability. *Random Structures & Algorithms*, 27(2):201-226, 2005.
- [15] Matthew Brennan and Guy Bresler. Reducibility and statistical-computational gaps from secret leakage. *Proceedings of 33rd COLT*, 648-847, 2020.
- [16] Matthew Brennan, Guy Bresler, Samuel B. Hopkins, Jerry Li, and Tselil Schramm. Statistical query algorithms and low-degree tests are almost equivalent. *Conference on Learning Theory (COLT)*, 2021.
- [17] Michael Celentano and Andrea Montanari. Fundamental barriers to high-dimensional regression with convex penalties. *arXiv:1903.10603*, preprint 2019.
- [18] Ming-Te Chao and John Franco. Probabilistic analysis of a generalization of the unit-clause literal selection heuristic for the k -satisfiability problem. *Information Sciences*, 51:289-314, 1990.
- [19] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Integrality Gaps for Sherali-Adams Relaxations. *Proceedings of 41st STOC*, 283-292, 2009.
- [20] Ruiwen Chen, Rahul Santhanam, and Srikanth Srinivasan. Average-case lower bounds and satisfiability algorithms for small threshold circuits. *Theory of Computing*, 14(9):1-55, 2018.
- [21] Wei-Kuo Chen, David Gamarnik, Dmitry Panchenko, and Mustazee Rahman. Suboptimality of local algorithms for a class of max-cut problems. *Annals of Probability*, 47(3):1587-1618, 2019.
- [22] Yeshwanth Cherapanamjeri, Samuel B. Hopkins, Tarun Kathuria, Prasad Raghavendra, and Nilesch Tripuraneni. Algorithms for heavy-tailed statistics: Regression, covariance estimation, and beyond. *Proceedings of 52nd STOC*, 601-609, 2020.
- [23] Václav Chvátal and Bruce Reed. Mick gets some (the odds are on his side). *Proceedings of 33th FOCS*, 620-627, 1992.
- [24] Václav Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759-768, 1988.
- [25] Amin Coja-Oghlan. A better algorithm for random k -SAT. *SIAM Journal on Computing*, 39:2823-2864, 2010.
- [26] Amin Coja-Oghlan and Charilaos Efthymiou. On independent sets in random graphs. *Random Structures & Algorithms*, 47(3):436-486, 2015.

- [27] Amin Coja-Oghlan, Uriel Feige, Alan Frieze, Michael Krivelevich, and Dan Vilenchik. On smoothed k -CNF formulas and the Walksat algorithm. *Proceedings of 20th SODA*, 451-460, 2009.
- [28] Amin Coja-Oghlan, Amir Haqshenas, and Samuel Hetterich. Walksat stalls well below the satisfiability threshold. *SIAM Journal on Discrete Mathematics*, 31:160-1173, 2017.
- [29] Amin Coja-Oghlan and Konstantinos Panagiotou. The asymptotic k -SAT threshold. *Advances in Mathematics*, 288:985-1068, 2016.
- [30] Stephen Cook. The complexity of theorem proving procedures. *Proceedings of 3rd STOC*, 151-158, 1971.
- [31] Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem proving. *Communications of the ACM*, 5(7):394-397, 1961.
- [32] Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201-205, 1960.
- [33] Jian Ding, Allan Sly, and Nike Sun. Proof of the satisfiability conjecture for large k . *Proceedings of 47th STOC*, 59-68, 2015.
- [34] Yunzi Ding, Dmitriy Kunisky, Alexander S Wein, and Afonso S. Bandeira. Subexponential-time algorithms for sparse PCA. *arXiv:1907.11635*, preprint 2019.
- [35] David L. Donoho, Arian Maleki, and Andrea Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914-18919, 2009.
- [36] Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. Optimization of mean-field spin glasses. *arXiv:2001.00904*, preprint 2020.
- [37] John Franco and Marvin Paull. Probabilistic analysis of the Davis–Putnam procedure for solving the satisfiability problem. *Discrete Applied Mathematics*, 5(1):77–87, 1983.
- [38] Alan Frieze. On the independence number of random graphs. *Discrete Mathematics*, 81(2):171-175, 1990.
- [39] Alan Frieze and Stephen Suen. Analysis of two simple heuristics on a random instance of k -SAT. *Journal of Algorithms*, 20:312-355, 1996.
- [40] Merrick L. Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13-27, 1984.
- [41] David Gamarnik and Aukosh Jagannath. The overlap gap property and approximate message passing algorithms for p -spin models. *Annals of Probability*, 2021.
- [42] David Gamarnik, Aukosh Jagannath, and Subhabrata Sen. The overlap gap property in principal sub-matrix recovery. *arXiv:1908.09959*, preprint 2019.
- [43] David Gamarnik, Aukosh Jagannath, and Alexander S. Wein. Low-degree hardness of random optimization problems. *Proceedings of 61st FOCS*, 2020.
- [44] David Gamarnik and Eren C. Kızıldağ. Algorithmic obstructions in the random number partitioning problem. *arXiv:2103.01369*, preprint 2021.
- [45] David Gamarnik and Madhu Sudan. Limits of local algorithms over sparse random graphs. *The Annals of Probability*, 45(4):2353-2376, 2017.
- [46] David Gamarnik and Madhu Sudan. Performance of sequential local algorithms for the random NAE- k -SAT problem. *SIAM Journal on Computing*, 46(2):590-619, 2017.
- [47] David Gamarnik and Ilias Zadik. The landscape of the planted clique problem: dense subgraphs and the overlap gap property. *arXiv:1904.07174*, preprint 2019.

- [48] Allen T. Goldberg, Paul W. Purdom, and Cynthia Brown. Average time analysis of simplified Davis-Putnam procedures. *Information Processing Letters*, 15:72–75, 1982.
- [49] Dima Grigoriev. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theoretical Computer Science*, 259(1-2):613-622, 2001.
- [50] Johan Håstad. Almost optimal lower bounds for small depth circuits. *Proceedings of 18th STOC*, 6-20, 1986.
- [51] Hamed Hatami, László Lovász, and Balázs Szegedy. Limits of local-global convergent graph sequences. *arXiv:1205.4356*, preprint 2012.
- [52] Samuel Hetterich. Analysing survey propagation guided decimation on random formulas. *Proceedings of 43rd ICALP*, #65, 2016
- [53] Samuel B. Hopkins. *Statistical Inference and the Sum of Squares Method*. PhD thesis, Cornell University, 2018.
- [54] Samuel B. Hopkins, Pravesh K. Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. The power of sum-of-squares for detecting hidden structures. *Proceedings of 58th FOCS*, 720-731, 2017.
- [55] Samuel B. Hopkins and David Steurer. Efficient bayesian estimation from few samples: community detection and related problems. *Proceedings of 58th FOCS*, 379-390, 2017.
- [56] Adel Javanmard and Andrea Montanari. State evolution for general approximate message passing algorithms, with applications to spatial coupling. *Information and Inference: A Journal of the IMA*, 2(2):115-144, 2013.
- [57] Richard M. Karp. The probabilistic analysis of some combinatorial search algorithms, in *Algorithms and Complexity: New Directions and Recent Results*, J. F. Traub, ed., Academic Press, New York, pp. 1-19, 1976.
- [58] Lefteris M. Kirousis, Evangelos Kranakis, Danny Krizanc, Yannis C. Stamatiou. Approximating the unsatisfiability threshold of random formulas. *Random Structures & Algorithms*, 12(3):253–269, 1998.
- [59] Pravesh K. Kothari, Ryuhei Mori, Ryan O’Donnell, and David Witmer. Sum of squares lower bounds for refuting any CSP. *Proceedings of 49th STOC*, 132-145, 2017.
- [60] Florent Krzakala, Andrea Montanari, Federico Ricci-Tersenghi, Guilhem Semerjian, and Lenka Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proceedings of the National Academy of Sciences*, 104:10318-10323, 2007.
- [61] Dmitriy Kunisky, Alexander S. Wein, and Afonso S. Bandeira. Notes on computational hardness of hypothesis testing: predictions using the low-degree likelihood ratio. *arXiv:1907.11636*, preprint 2019.
- [62] Michael G. Luby, Michael Mitzenmacher, and M. Amin Shokrollahi. Analysis of random processes via and-or tree evaluation. *Proceedings of 9th SODA*, 364-373, 1998.
- [63] Yuetian Luo and Anru R. Zhang. Tensor clustering with planted structures: statistical optimality and computational limits. *arXiv:2005.10743*, 2020.
- [64] Andrea Montanari. Optimization of the Sherrington-Kirkpatrick hamiltonian. *Proceedings of 60th FOCS*, 1417-1433, 2019.
- [65] Andrea Montanari, Federico Ricci-Tersenghi, and Guilhem Semerjian. Solving constraint satisfaction problems through belief propagation-guided decimation. *Proceedings of 45th Allerton*, 352-359, 2007.
- [66] Danny Nam, Allan Sly, and Youngtak Sohn. One-step replica symmetry breaking of random regular NAE- k -SAT. *arXiv:2011.14270*, preprint 2020.

- [67] Ryan O'Donnell. *Analysis of Boolean functions*. Cambridge University Press, 2014.
- [68] Christos H. Papadimitriou. On selecting a satisfying truth assignment. *Proceedings of 32 FOCS*, 163-169, 1991.
- [69] Mustazee Rahman and Bálint Virág. Local algorithms for independent sets are half-optimal. *The Annals of Probability*, 45(3):1543-1577, 2017.
- [70] Tselil Schramm and Alexander S. Wein. Computational barriers to estimation from low-degree polynomials. *arXiv:2008.02269*, preprint 2020.
- [71] Alexander S. Wein. Optimal low-degree hardness of maximum independent set. *arXiv:2010.06563*, preprint 2020.

A Small Hamming Distance Implies Small Conditional Overlap Entropy

In this section, we present the deferred proof of Lemma 4.8, which shows that a small change in $x \in \{\mathbf{T}, \mathbf{F}\}^n$ causes only a small change in $H(\pi(x|y^{(0)}, \dots, y^{(\ell-1)}))$.

Lemma 4.8. *Let $\ell \in \mathbb{N}$ be arbitrary and let $x, x', y^{(0)}, \dots, y^{(\ell-1)} \in \{\mathbf{T}, \mathbf{F}\}^n$. If $\Delta(x, x') \leq \frac{1}{2}$, then*

$$\left| H\left(\pi(x|y^{(0)}, \dots, y^{(\ell-1)})\right) - H\left(\pi(x'|y^{(0)}, \dots, y^{(\ell-1)})\right) \right| \leq H(\Delta(x, x')).$$

The $H(\cdot)$ on the right denotes the binary entropy function.

Proof. For each partition $\{S, T\} \in \mathcal{P}_2(\ell)$, let

$$I_{S,T} = \left\{ i \in [n] : \text{all } \{y_i^{(t)} : t \in S\} \text{ equal one value and all } \{y_i^{(t)} : t \in T\} \text{ equal the other value} \right\}.$$

Note that $|I_{S,T}| = \pi_{S,T}n$. If $\pi_{S,T} \neq 0$, define

$$\lambda_{S,T} = \frac{1}{|I_{S,T}|} \#(i \in I_{S,T} : x_i = \mathbf{T}) \quad \text{and} \quad \lambda'_{S,T} = \frac{1}{|I_{S,T}|} \#(i \in I_{S,T} : x'_i = \mathbf{T}).$$

(If $\pi_{S,T} = 0$, we can set these values arbitrarily in $[0, 1]$.) Thus,

$$\frac{1}{2} \geq \Delta(x, x') \geq \sum_{\{S,T\} \in \mathcal{P}_2(\ell)} \pi_{S,T} |\lambda_{S,T} - \lambda'_{S,T}|.$$

Let $\sum_{\{S,T\} \in \mathcal{P}_2(\ell)} \pi_{S,T} |\lambda_{S,T} - \lambda'_{S,T}| = \mu$. Thus $\mu \leq \Delta(x, x') \leq \frac{1}{2}$. Moreover, from the definition of conditional overlap entropy,

$$H\left(\pi(x|y^{(0)}, \dots, y^{(\ell-1)})\right) = \sum_{\{S,T\} \in \mathcal{P}_2(\ell)} \pi_{S,T} H(\lambda_{S,T}),$$

and analogously for x' . Note that $H(\cdot)$ is concave, so $H'(\cdot)$ is decreasing. Thus, for all $[a, b] \in [0, 1]$ with $a \geq b$,

$$H(a) - H(b) = \int_a^b H'(x) \, dx \leq \int_0^{a-b} H'(x) \, dx = H(a-b).$$

Similarly $H(1-b) - H(1-a) \leq H(a-b)$, whence $|H(a) - H(b)| \leq H(a-b)$. Thus,

$$\begin{aligned} \left| H\left(\pi(x|y^{(0)}, \dots, y^{(\ell-1)})\right) - H\left(\pi(x'|y^{(0)}, \dots, y^{(\ell-1)})\right) \right| &\leq \sum_{\{S,T\} \in \mathcal{P}_2(\ell)} \pi_{S,T} |H(\lambda_{S,T}) - H(\lambda'_{S,T})| \\ &\leq \sum_{\{S,T\} \in \mathcal{P}_2(\ell)} \pi_{S,T} H(|\lambda_{S,T} - \lambda'_{S,T}|). \end{aligned}$$

By concavity of $H(\cdot)$, this last quantity has maximum value $H(\mu)$, attained when all the $|\lambda_{S,T} - \lambda'_{S,T}|$ are equal to μ . Because $H(\cdot)$ is increasing on $[0, \frac{1}{2}]$ and $\mu \leq \Delta(x, x') \leq \frac{1}{2}$, we conclude that

$$\left| H\left(\pi(x|y^{(0)}, \dots, y^{(\ell-1)})\right) - H\left(\pi(x'|y^{(0)}, \dots, y^{(\ell-1)})\right) \right| \leq H(\mu) \leq H(\Delta(x, x')).$$

□

B On Improving the Constant κ^*

In this section, we discuss how the constant κ^* in Theorem 2.6 can be improved. We define a constant κ^{**} as the solution to a maximin problem. We will show that $\kappa^{**} \leq \kappa^*$ and sketch how our proof of Theorem 2.6 can be lightly modified to improve the constant κ^* to κ^{**} . We then heuristically argue that $\kappa^{**} < \kappa^*$, so that this modification is an improvement. We also prove that κ^{**} is bounded below by a constant larger than 1, approximately 1.716, and thus further ideas are needed to fully close the constant-factor gap between Theorems 2.6 and 2.10. Because κ^{**} remains bounded away from 1, and we believe 1 is the optimal constant, we did not attempt to rigorously evaluate κ^{**} . The full arguments given for κ^* in the body of the paper are also more intuitive than those for κ^{**} .

B.1 A Maximin Problem

Let (Ξ, P_ξ) be an arbitrary probability space. Let $\mathcal{Q}$ be the space of functions $q : \Xi \rightarrow [0, 1]$, equipped with the metric $d(q, q') = \mathbb{E}_\xi |q(\xi) - q'(\xi)|$. For $q \in \mathcal{Q}$, let $D(q)$ be the law of u sampled by the following experiment: sample $\xi \sim (\Xi, P_\xi)$; then, with probability $q(\xi)$, set $u = -\log q(\xi)$ and otherwise set $u = -\log(1 - q(\xi))$. Clearly $\mathbb{E}_{u \sim D(q)} u = \mathbb{E}_\xi H(q(\xi))$. Define

$$F(q) = \frac{1}{\log k} \cdot \frac{\log 2 + k \mathbb{E}_\xi H(q(\xi))}{\mathbb{P}_{(u_1, \dots, u_k) \sim D(q)^{\otimes k}} \left[\sum_{i=1}^k u_i \geq \log k + \log \log k \right]}.$$

Let $\mathcal{P}$ be the set of functions $p : \Xi \times [0, 1] \rightarrow [0, 1]$, such that $p(\xi, 0) \in \{0, 1\}$ and $p(\xi, 1) = \frac{1}{2}$ for all $\xi \in \Xi$, and $p(\cdot, s)$ (which, for fixed $s \in [0, 1]$, is an element of $\mathcal{Q}$) is continuous in s with respect to the topology of $\mathcal{Q}$. Consider the maximin problem

$$\kappa^{**} = \limsup_{k \rightarrow \infty} \max_{p \in \mathcal{P}} \min_{s \in [0, 1]} F(p(\cdot, s)). \quad (\text{B.1})$$

This has the following geometric interpretation: κ^{**} is the smallest constant such that the sub-level set $\{q \in \mathcal{Q} : F(q) \leq \kappa^{**}\}$ topologically disconnects the functions $q \equiv 0$ and $q \equiv \frac{1}{2}$ in $\mathcal{Q}$. (Note that $\mathcal{Q}$ is symmetric under replacing $q(\xi)$ with $1 - q(\xi)$ for any subset of the $\xi \in \Xi$, and $F(q) = F(q')$ for any q, q' related by such a symmetry. Thus, equivalently κ^{**} is the smallest constant such that this sub-level set disconnects the function $q \equiv \frac{1}{2}$ from any $q \in \mathcal{Q}$ with $q(\xi) \in \{0, 1\}$ for all $\xi \in \Xi$.)

First, we show that κ^* is an upper bound on the solution to this maximin problem.

Proposition B.1. *We have that $\kappa^* \geq \kappa^{**}$.*

Proof. Fix some $p \in \mathcal{P}$. By continuity of $p(\cdot, s)$ in s , we can set $s \in [0, 1]$ such that $\mathbb{E}_\xi H(p(\xi, s)) = \beta^* \frac{\log k}{k}$. As in the proof of Proposition 5.7, we apply a Chernoff bound on the random variables $\frac{\min(u_i, \log k)}{\log k}$ to show that, for any $\beta > 1$ and $q : \Omega \rightarrow [0, 1]$ with $\mathbb{E}_\xi H(q(\xi)) = \beta \frac{\log k}{k}$, we have

$$\mathbb{P}_{(u_1, \dots, u_k) \sim D(q)^{\otimes k}} \left[\sum_{i=1}^k u_i < \log k + \log \log k \right] \leq \beta e^{-(\beta-1)} + o_k(1). \quad (\text{B.2})$$

In particular, for the s we chose,

$$F(p(\cdot, s)) \leq \frac{\frac{2}{\log k} + \beta^*}{1 - \beta^* e^{-(\beta^*-1)} - o_k(1)} \rightarrow \iota(\beta^*) = \kappa^*.$$

□

Next, we sketch how the proof of Theorem 2.6 can be improved to replace κ^* with κ^{**} .

Proposition B.2. *Theorem 2.6 holds for all $\kappa > \kappa^{**}$.*

Proof Sketch. Like in the original proof of Theorem 2.6, we let $x^{(t)} = \mathcal{A}(\Phi^{(t)})$ for $0 \leq t \leq T$. On the event S_{valid} and analogues of S_{consec} , S_{indep} , S_{ogp} , we will pick $k+1$ indices $0 \leq t_0 \leq t_1 \leq \dots \leq t_k \leq T$ and assignments $y^{(\ell)} = x^{(t_\ell)}$ violating the analogue of S_{ogp} , yielding a contradiction.

We set $t_0 = 0$. For $1 \leq \ell \leq k$ we set t_ℓ to be the smallest $t > t_{\ell-1}$ obeying a stopping rule in terms of $x^{(t)}, y^{(0)}, \dots, y^{(\ell-1)}$, which we show occurs at or before time $t = t_{\ell-1} + km$. We will describe the stopping rule explicitly for $y^{(0)} = \mathbf{T}^n$. In general, the stopping rule holds for $x^{(t)}, y^{(0)}, y^{(1)}, \dots, y^{(\ell-1)}$ if it holds for $x^{(t)} \oplus y^{(0)}, \mathbf{T}^n, y^{(1)} \oplus y^{(0)}, \dots, y^{(\ell-1)} \oplus y^{(0)}$, where $\oplus$ denotes bitwise XOR.

We now fix some $y^{(0)}, \dots, y^{(\ell-1)}$, with $y^{(0)} = \mathbf{T}^n$. The probability space Ξ abstracts the sample space of $y_i^{(\leq \ell-1)}$, where $i \in \text{unif}([n])$. For $t_{\ell-1} \leq t \leq t_{\ell-1} + km$, let $\phi_{\ell,t}(y_i^{(\leq \ell-1)})$ denote $\phi_\ell(\mathbf{T}|y_i^{(\leq \ell-1)})$ if we set $y^{(\ell)} = x^{(t)}$. We construct $p \in \mathcal{P}$ as follows. For $s \in [0, \frac{1}{2}]$, we let $p(\cdot, s)$ interpolate continuously through the functions $\phi_{\ell,t}$ from $t = t_{\ell-1}$ (at $s = 0$) to $t = t_{\ell-1} + km$ (at $s = \frac{1}{2}$). Using the analogue of S_{indep} , we have that $\mathbb{E}_\xi H(p(\xi, \frac{1}{2})) \geq (\kappa - \varepsilon) \frac{\log k}{k}$ for a constant $\varepsilon > 0$, small enough that $\kappa - \varepsilon > \kappa^*$. For $s \in [\frac{1}{2}, 1]$, we let $p(\cdot, s)$ evolve continuously to $p(\cdot, 1) \equiv \frac{1}{2}$, such that for all $s \in [\frac{1}{2}, 1]$ we have $\mathbb{E}_\xi H(p(\xi, s)) \geq (\kappa - \varepsilon) \frac{\log k}{k}$. Thus we have $F(p(\cdot, s)) \geq \kappa - \varepsilon$ for all $s \in [\frac{1}{2}, 1]$. Because $F(p(\cdot, s)) \leq \kappa^{**}$ for some s for every $p \in \mathcal{P}$, the minimum of $F(p(\cdot, s))$ occurs on $s \in [0, \frac{1}{2}]$. Using the analogue of S_{consec} , we have that one of the functions $\phi_{\ell,t}$ occurs near this minimum, in the sense that $F(\phi_{\ell,t}) \leq \kappa - \varepsilon$. The stopping rule chooses t_ℓ to be this t .

The analogue of S_{ogp} is the event that no $0 \leq t_0 \leq t_1 \leq \dots \leq t_k \leq T$ and $y^{(0)}, \dots, y^{(k)} \in \{\mathbf{T}, \mathbf{F}\}^n$ exist where $y^{(\ell)}$ satisfies $\Phi^{(t_\ell)}$ for all $0 \leq \ell \leq k$ and $y^{(0)}, \dots, y^{(k)}$ have the overlap profile implicitly described above. We can show $\mathbb{P}(S_{\text{ogp}}^c) \leq \exp(-\Omega(n))$ by techniques analogous to Section 5. The key point is that the arguments in that section require precisely that $F(q)$ is bounded below κ for each $q = \phi_\ell(\mathbf{T}|\cdot)$. This yields the desired contradiction. $\square$

B.2 Suboptimality of κ^*

We believe that $\kappa^* > \kappa^{**}$ due to the following heuristic argument. The Chernoff bound (B.2) is tightest when most of the mass of the random variables $\frac{\min(u_i, \log k)}{\log k}$ is near 0 or 1. When this occurs, most of the mass of u_i is near 0 or $\log k$. Then, the event that $\sum_{i=1}^k u_i \geq \log k + \log \log k$ is the event that one or two of the u_i attains a value near $\log k$. This is a tail probability in a non-asymptotic regime – approximately, the probability that a Poisson random variable is larger than 1 or 2 – so the Chernoff bound will not get exactly the correct probability.

B.3 κ^{**} is Bounded Away from 1

In this section, we will show that κ^{**} is bounded below by a constant, approximately 1.716, larger than 1. Thus, a generic application of our methods cannot fully close the constant-factor gap between Theorems 2.6 and 2.10.

We will first show a weaker lower bound on κ^{**} . Define $\psi_1 : (0, +\infty) \rightarrow \mathbb{R}$ by

$$\psi_1(\lambda) = \frac{\lambda/2}{1 - (1 + \lambda)e^{-\lambda}},$$

and let $\psi_1^* = \min_{\lambda > 0} \psi_1(\lambda) \approx 1.675$.

Proposition B.3. *We have $\kappa^{**} \geq \psi_1^*$.*

Proof. We will prove this proposition by constructing a suitable function family $p \in \mathcal{P}$.

Let $\Xi = [0, 1]$ equipped with the uniform measure. Let $p : \Xi \times [0, 1] \rightarrow [0, 1]$ be defined by

$$p(\xi, s) = \begin{cases} \min(s, \frac{1}{2}) & \xi \leq s, \\ 0 & \xi \geq s. \end{cases}$$

Thus, for fixed $s \in [0, 1]$, $p(\xi, s) = \min(s, \frac{1}{2})$ with probability s , and otherwise $p(\xi, s) = 0$. We will show that for this p ,

$$\limsup_{k \rightarrow \infty} \min_{s \in [0, 1]} F(p(\cdot, s)) \geq \psi_1^*,$$

from which the proposition follows.

Note that if $s = \omega_k(k^{-1/2})$, then $\mathbb{E}_\xi H(p(\xi, s)) = \omega_k(\log k/k)$, and so $F(p(\cdot, s)) = \omega_k(1)$. Therefore it suffices to consider $s = O_k(k^{-1/2})$. Then,

$$\mathbb{E}_\xi H(p(\xi, s)) = (1 + o_k(1))s^2 \log \frac{1}{s}.$$

We now analyze the behavior of the denominator of $F(p(\cdot, s))$. Note that a sample $u \sim D(p(\cdot, s))$ equals $\log \frac{1}{s}$ with probability s^2 , $\log \frac{1}{1-s} \leq \frac{s}{1-s}$ with probability $s(1-s)$, and 0 with probability $1-s$. For $i = 1, \dots, k$, define

$$v_i = \log \frac{1}{s} \mathbb{1} \left\{ u_i = \log \frac{1}{s} \right\}, \quad \text{and} \quad w_i = \frac{s}{1-s} \mathbb{1} \left\{ u_i = \log \frac{1}{1-s} \right\}.$$

So, $u_i \leq v_i + w_i$. For $u_1, \dots, u_k \sim D(p(\cdot, s))^{\otimes k}$, we have

$$\mathbb{P} \left[\sum_{i=1}^k u_i \geq \log k + \log \log k \right] \leq \mathbb{P} \left[\sum_{i=1}^k v_i \geq \log k \right] + \mathbb{P} \left[\sum_{i=1}^k w_i \geq \log \log k \right].$$

Let $1 + \delta = \frac{\log \log k}{k \mathbb{E} w_1} = \frac{\log \log k}{k s^2}$. Because $s = O_k(k^{-1/2})$, we have $\frac{\delta^2}{2+\delta} \geq \frac{1}{2}(1 + \delta)$. By a Chernoff bound,

$$\begin{aligned} \mathbb{P} \left[\sum_{i=1}^k w_i \geq \log \log k \right] &\leq \mathbb{P} \left[\sum_{i=1}^k \frac{1-s}{s} w_i \geq \frac{1-s}{s} \log \log k \right] \leq \exp \left(-\frac{\delta^2}{2+\delta} \cdot k s (1-s) \right) \\ &\leq \exp \left(-\frac{1}{2}(1+\delta) k s (1-s) \right) \leq \exp \left(-\frac{(1-s) \log \log k}{2s} \right) \leq \exp \left(-\Omega_k(k^{-1/2}) \right). \end{aligned}$$

To analyze the other probability, we consider cases $s > \frac{1}{k}$ and $s \leq \frac{1}{k}$. We first consider $s > \frac{1}{k}$. In order to have $\sum_{i=1}^k v_i \geq \log k$, at least two v_i must be nonzero. This occurs with probability

$$1 - (1 - s^2)^k - s^2 k (1 - s^2)^{k-1} \leq 1 - (1 + s^2 k)(1 - s^2)^k.$$

Thus,

$$F(p(\cdot, s)) \geq \frac{1}{\log k} \cdot \frac{\log 2 + (1 + o_k(1))s^2 k \log \frac{1}{s}}{1 - (1 + s^2 k)(1 - s^2)^k + \exp(-\Omega_k(k^{-1/2}))}.$$

If $s^2 k = o_k(1)$, then $1 - (1 + s^2 k)(1 - s^2)^k = O_k(s^4 k^2)$, and the right-hand side is $\omega_k(1)$. So, this bound is minimized at $s = \lambda k^{-1/2}$ for constant λ , in which case

$$\frac{1}{\log k} \cdot \frac{\log 2 + (1 + o_k(1))s^2 k \log \frac{1}{s}}{1 - (1 + s^2 k)(1 - s^2)^k + \exp(-\Omega_k(k^{-1/2}))} \rightarrow \frac{\lambda/2}{1 - (1 + \lambda) \exp(-\lambda)} = \psi_1(\lambda) \geq \psi_1^*.$$

We now consider $s \leq \frac{1}{k}$. In order to have $\sum_{i=1}^k v_i \geq \log k$, at least one v_i must be nonzero. This occurs with probability

$$1 - (1 - s^2)^k = (1 + o_k(1))s^2 k,$$

and so

$$F(p(\cdot, s)) \geq \frac{1}{\log k} \cdot \frac{\log 2 + (1 + o_k(1))s^2 k \log \frac{1}{s}}{(1 + o_k(1))s^2 k + \exp(-\Omega_k(k^{-1/2}))}.$$

The right-hand side is $\omega_k(1)$ because $s \leq \frac{1}{k}$. □

For any nonnegative integer N , we may further define

$$\psi_N(\lambda) = \frac{\lambda/(N+1)}{\left(\sum_{i=0}^N \frac{\lambda^i}{i!}\right) \exp(-\lambda)}$$

and $\psi_N^* = \inf_{\lambda>0} \psi_N(\lambda)$. Over positive integers N , the largest ψ_N^* is $\psi_2^* \approx 1.716$. The following corollary gives the lower bound on κ^{**} alluded to above.

Corollary B.4. *We have that $\kappa^{**} \geq \psi_2^*$.*

Proof. We will construct a suitable function family p . For any nonnegative integer N , we can define

$$p_N(\xi, s) = \begin{cases} \min(s, \frac{1}{2}) & \xi \leq s^N, \\ 0 & \xi \geq s. \end{cases} \quad (\text{B.3})$$

By a similar analysis to Proposition B.3, we can show for this p that

$$\limsup_{k \rightarrow \infty} \min_{s \in [0,1]} F(p_N(\cdot, s)) \geq \psi_N^*.$$

Taking $N = 2$ yields the result. □

Due to Corollary B.4, a proof along the lines of the proof of Theorem 2.6 that improves the constant κ^* below ψ_2^* must pick the sequence $y^{(0)}, \dots, y^{(\ell)}$ in a more sophisticated way. One such improvement may be to pick the sequence $y^{(0)}, \dots, y^{(k)}$ all at once and after seeing the entire sequence $x^{(0)}, \dots, x^{(T)}$, rather than by revealing the $x^{(t)}$ one by one and picking the $y^{(\ell)}$ by a stopping rule. Another such improvement may be to extract further properties of the sequence $x^{(0)}, \dots, x^{(T)}$; we currently only use that this sequence is stable in the sense of S_{consec} .

We conjecture that Corollary B.4 is in fact sharp.

Conjecture B.5. *We have that $\kappa^{**} = \psi_2^*$. In particular, Theorem 2.6 holds for all $\kappa > \psi_2^*$.*

The following evidence supports this conjecture. In the maximin problem (B.1), if we restrict the maximum over p to “two-layer” p – that is, p such that for every s , $p(\xi, s)$ attains at most one nonzero value – then we can show by explicit computation that the maximin problem has value κ^{**} . The key idea of this proof is that for each two-layer p , at the s minimizing $F(p(\cdot, s))$, $p(\cdot, s)$ behaves like $p_N(\cdot, s')$ for some s' and some (possibly fractional) N . We can show that taking N to be fractional does not maximize $\min_{s \in [0,1]} F(p_N(\cdot, s))$. Thus the candidate maxima are p_N for integer N , and of these p_2 is maximal, attaining value ψ_2^* . We believe that the maximum of (B.1) over $p \in \mathcal{P}$ is attained by a two-layer p .