The Asymptotic k -SAT Threshold

[Extended Abstract]

Amin Coja-Oghlan^{*}
Goethe University
Mathematics Institute
Frankfurt 60054, Germany
acoghlan@math.uni-frankfurt.de

ABSTRACT

Since the early 2000s physicists have developed an ingenious but non-rigorous formalism called the *cavity method* to put forward precise conjectures as to the phase transitions in random constraint satisfaction problems (“CSPs”). The cavity method comes in two versions: the simpler *replica symmetric* variant, and the more intricate *1-step replica symmetry breaking* (“1RSB”) version. While typically the former only gives upper and lower bounds, the latter is conjectured to yield precise results in many cases. By now, there are a number of examples where the replica symmetric bounds have been verified rigorously. However, verifications of 1RSB predictions are scarce. Perhaps the most prominent challenge in this context is that of pinning down the random k -SAT threshold $r_{k\text{-SAT}}$. Here we prove that $r_{k\text{-SAT}} = 2^k \ln 2 - \frac{1}{2}(1 + \ln 2) + o_k(1)$, which matches the 1RSB prediction up to the $o_k(1)$ error term. The proof directly employs ideas from the 1RSB cavity method, such as the notion of covers (relaxed satisfying assignments) and bits of the Survey Propagation calculations. The best previous lower bound was $r_{k\text{-SAT}} \geq 2^k \ln 2 - \frac{3}{2} \ln 2 + o_k(1)$, matching the replica symmetric lower bound asymptotically [Coja-Oghlan, Panagiotou: STOC 2013].

Categories and Subject Descriptors

G.2 [Discrete Mathematics]: General; F.2 [Analysis of Algorithms and Problem Complexity]: General

General Terms

Theory

^{*}The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 278857-PTCC.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s)
STOC’14, May 31–June 3, 2014, New York, NY, USA.
ACM 978-1-4503-2710-7/14/05.
<http://dx.doi.org/10.1145/2591796.2591822>.

Keywords

Random Structures, Phase Transitions, k -SAT, Second Moment Method, Survey Propagation

1. INTRODUCTION

Throughout the paper, we let $k \geq 3$ and n denote integers and we let $V = \{x_1, \dots, x_n\}$ be a set of variables taking values 0 (‘false’) or 1 (‘true’). Moreover, $\Phi = \Phi_k(n, m)$ signifies a random k -CNF formula over the variables V with m clauses. We generally assume that $m = \lceil rn \rceil$ for a fixed number $r > 0$, the *density*.

Since the early 1990s experimental work has been supporting two hypotheses about the random k -SAT problem [6, 23]. First, that there is a *threshold for satisfiability*, i.e., for any $k \geq 3$ there is a critical density $r_{k\text{-SAT}}$ where the probability that the random formula is satisfiable drops from nearly 1 to nearly 0. Second, that for densities close to but below this threshold finding a satisfying assignment is algorithmically challenging. An impressive bulk of theoretical work has since augmented this picture. Yet in spite of over twenty years of research, random formulas near the satisfiability threshold retain a seasoned reputation for being computationally difficult. And pinning down the as yet elusive satisfiability threshold has become *the* benchmark problem in the study of random constraint satisfaction problems.

From its early days the random k -SAT problem has been studied via methods from statistical physics. Through the physics lens, random k -SAT is an example of a “disordered system”, and physicists have developed a systematic albeit non-rigorous approach to this type of problem called the *cavity method* [31]. In fact, the cavity method comes in two installments. The simpler version is the so-called *replica symmetric ansatz*, associated with the *Belief Propagation* message passing algorithm. Unfortunately, the replica symmetric variant does not suffice to calculate the k -SAT threshold precisely. It merely predicts upper and lower bounds [34]: expanded asymptotically in k ,

$$r_{\text{cond}} = 2^k \ln 2 - \frac{3 \ln 2}{2} - o_k(1) \leq r_{k\text{-SAT}} \leq 2^k \ln 2 - \frac{\ln 2}{2}. \quad (1)$$

Here and throughout, $o_k(1)$ signifies a term that tends to 0 in the limit of large k . Moreover, the Belief Propagation algorithm provably fails to find satisfying assignments well below the k -SAT threshold [9].

Hence the need for the second, far more powerful but also far more intricate version, the *1-step replica symmetry breaking ansatz* (“1RSB”), and its associated *Survey Propagation*

message passing procedure [32]. The 1RSB cavity method is conjectured to predict the precise thresholds not only in random k -SAT, but also in a host of other problems ranging from classical physics models to low-density parity check codes to compressed sensing. Specifically, the 1RSB prediction on the k -SAT threshold, expanded asymptotically in terms of k , reads [30]

$$r_{k\text{-SAT}} = 2^k \ln 2 - \frac{1 + \ln 2}{2} + o_k(1). \quad (2)$$

Furthermore, the Survey Propagation algorithm currently is the most successful algorithm experimentally for solving random k -SAT instances [5, 25].¹

While the numerical difference between the bounds from (1) and the prediction (2) may seem modest, there is a substantial conceptual difference between the replica symmetric and the 1RSB cavity method. In fact, the lower bound in (1) is expected to mark a phase transition called *condensation*, at which the probabilistic nature of the problem changes dramatically [26]. The 1RSB cavity method is designed to work beyond this point.

From the viewpoint of the cavity method as well as from a rigorous perspective, random k -SAT is by far the most challenging problem among the standard examples of random CSPs that have been studied intensively. The reason is that in satisfiability there is a fundamental asymmetry between the two “spins”, i.e., the Boolean values ‘true’ and ‘false’. More specifically, consider the (thought) experiment of first generating a random formula Φ and then sampling a random satisfying assignment σ of Φ . Then the local “shape” of Φ holds significant clues as to the probability that a given variable x takes the value ‘true’ under the random assignment σ . For instance, if x appears many more times positively than negatively in Φ , then we should expect that the probability that x takes the value ‘true’ under σ is greater than $1/2$. This asymmetry, which is provably present [11], is in contrast to, e.g., the graph coloring problem, where all the colors have the same “meaning”. In fact, the probability that a given vertex takes a particular color in a random coloring is just uniform, simply because we can permute the color classes. Similarly, the k -NAESAT (“Not-All-Equal-Satisfiability”) problem, which asks for a satisfying assignment whose inverse assignment is also satisfying, is perfectly symmetric by its very definition.²

In the present paper we develop a novel rigorous approach that allows us to prove the 1RSB cavity prediction (2). Coping with the asymmetry of the k -SAT problem will be a key challenge in the course of this. But before we state the main result, let us summarize the state of the art prior to the present work.

Friedgut [18] proved the existence of a (non-uniform) sat-

¹The author is unaware of any (non-trivial) analytical bounds on the density up to which the Survey Propagation algorithm finds satisfying assignments in random k -SAT. Generally, understanding the algorithms’ performance accurately appears to be even more difficult than calculating the k -SAT threshold. In fact, solving the latter problem might well be a precondition for tackling the former, and this is one motivation for the current work.

²Formally, we could call a random CSP *symmetric* if in a random problem instance for each variable the marginal distribution over the possible values that the variable can take (‘true’ or ‘false’ in satisfiability; the colors in graph coloring, etc.) converges to the uniform distribution.

isfiability threshold $r_{k\text{-SAT}}$ for any $k \geq 3$.³ However, his approach only establishes existence, without revealing the location of the k -SAT threshold. A “first moment” calculation yields an upper bound on the k -SAT threshold, namely $r_{k\text{-SAT}} \leq 2^k \ln 2 - \ln 2/2 + o_k(1)$. This matches the upper bound (1) that the replica symmetric cavity method predicts. Furthermore, a more sophisticated first moment argument [24] actually shows an upper bound of

$$r_{k\text{-SAT}} \leq 2^k \ln 2 - \frac{1}{2}(1 + \ln 2) + o_k(1),$$

which coincides with the 1RSB prediction (2).

Establishing *lower* bounds on $r_{k\text{-SAT}}$ is even more challenging. Proving $r_{k\text{-SAT}} \geq 2^k \ln 2 - \ln 2 - o_k(1)$, in [11] we turned the replica symmetric prediction (1) into a rigorous lower bound, the best prior one for general (large) values of k . The proof is based on the second moment method, whose use in the context of random k -SAT was pioneered by Achlioptas and Moore [2], Frieze and Wormald [20] and Achlioptas and Peres [3].

The main technical contribution of the present paper is a more powerful “1RSB version” of the second moment method. While the second moment argument in [11] was based on the (replica symmetric) Belief Propagation message passing procedure, the present argument uses the mightier Survey Propagation method and the associated notion of “relaxed” satisfying assignments called *covers*. Based on this novel approach, we establish

THEOREM 1. *The k -SAT threshold satisfies*

$$r_{k\text{-SAT}} = 2^k \ln 2 - \frac{1 + \ln 2}{2} + o_k(1).$$

There are several rigorous results that verify predictions based on the replica symmetric cavity method (see Section 3 below). But Theorem 1 is among the first rigorous theorems to verify a 1RSB prediction (at least up to the $o_k(1)$ error term), and it is the very first such rigorous result in an asymmetric problem. In order to cope with asymmetry, the proof of Theorem 1 explicitly incorporates the Survey Propagation calculations. To explain this approach, we need to take a closer look at the physics intuition.

2. SURVEY PROPAGATION, COVERS AND THE SECOND MOMENT METHOD

2.1 The second moment method

As pointed out in the seminal paper by Achlioptas and Moore [2], the second moment method can be used to prove lower bounds on the k -SAT threshold. The general strategy is as follows. Suppose that $Y = Y(\Phi) \geq 0$ is a random variable such that $Y(\Phi) > 0$ only if Φ is satisfiable. Assume, moreover, that there is a number $C = C(k) > 0$ that may depend on k but not on n such that

$$0 < \mathbb{E}[Y^2] \leq C \cdot \mathbb{E}[Y]^2. \quad (3)$$

Then the *Paley-Zygmund inequality*

$$\mathbb{P}[Y > 0] \geq \mathbb{E}[Y]^2 / \mathbb{E}[Y^2]$$

³Formally, for any $k \geq 3$ there is a sequence $r_{k\text{-SAT}}(n)$ such that for any fixed $\varepsilon > 0$, Φ is satisfiable w.h.p. if $m/n < (1 - \varepsilon)r_{k\text{-SAT}}(n)$ and unsatisfiable w.h.p. if $m/n > (1 + \varepsilon)r_{k\text{-SAT}}(n)$. The sequence $r_{k\text{-SAT}}(n)$ is conjectured to converge.

implies that

$$\liminf_{n \rightarrow \infty} \mathbb{P}[\Phi \text{ is sat}] \geq \liminf_{n \rightarrow \infty} \mathbb{P}[Y > 0] \geq 1/C > 0. \quad (4)$$

The following consequence of Friedgut’s sharp threshold theorem turns (4) into a lower bound on $r_{k\text{-SAT}}$.

LEMMA 1 ([18]). *If for some density r we have*

$$\liminf_{n \rightarrow \infty} \mathbb{P}[\Phi \text{ is satisfiable}] > 0,$$

then $r_{k\text{-SAT}} \geq r - o(1)$.

Thus, we “just” need to come up with a random variable Y that satisfies (3).

2.2 The replica symmetric ansatz

The obvious candidate seems to be the number Z of satisfying assignments of Φ . Of course, the second moment $\mathbb{E}[Z^2]$ is nothing but the expected number of *pairs* of satisfying assignments. In effect, it turns out that a necessary condition for the success of the second moment method is that in a *random pair* (σ, τ) of satisfying assignments of Φ , σ, τ “look uncorrelated”. A precise formal statement is given in [3, Section 2.3], but for instance, (3) can only hold if the average Hamming distance of σ, τ is $(1 + o(1))\frac{n}{2}$. However, in random k -SAT this is simply not the case. In effect, the bound (3) does *not* hold for Z for any $r > 0$.

As observed in [2, 3], the source of these correlations is the asymmetry of the k -SAT problem. More precisely, let d_{x_i} denote the *degree* of the variable x_i , i.e., number of times that x_i occurs positively in the formula Φ , and let $d_{\neg x_i}$ be the degree of $\neg x_i$. Furthermore, consider the *majority vote* assignment σ_{maj} , where we let $\sigma_{\text{maj}}(x_i) = 1$ if $d_{x_i} \geq d_{\neg x_i}$ and $\sigma_{\text{maj}}(x_i) = 0$ otherwise. Clearly, if the only information that we are given about Φ is the literal degrees $d_{x_1}, d_{\neg x_1}, \dots, d_{x_n}, d_{\neg x_n}$, then σ_{maj} is the assignment with the greatest probability of being satisfying. This is because σ_{maj} maximizes the number of true literal occurrences throughout the formula. To be precise, out of the km literals a

$$w_{\text{maj}} = \frac{1}{km} \sum_{i=1}^n \max\{d_{x_i}, d_{\neg x_i}\}$$

fraction are satisfied under σ_{maj} . Indeed, the closer an assignment σ is to σ_{maj} in Hamming distance, the larger the expected number of true literal occurrences. As a consequence, we expect that most satisfying assignments “lean towards” the majority assignment σ_{maj} . This induces a subtle correlation amongst pairs of satisfying assignments, which dooms the second moment method.

Nonetheless, as demonstrated in [11], by moving to a modified probability space it is possible to control the correlations induced by the drift towards σ_{maj} . The construction involves conditioning on the degrees $d_{x_i}, d_{\neg x_i}$ and designing a skewed probability distribution over assignments in which “uncorrelated” means that two assignments have exactly the correct drift towards σ_{maj} . In order to “guess” this skewed probability distribution, we used the Belief Propagation calculations from physics, thereby establishing the replica symmetric lower bound (1) on $r_{k\text{-SAT}}$ rigorously.

2.3 Condensation and Survey Propagation

However, this approach has no chance of actually yielding (2). The reason is that, according to the cavity method,

at the density $r_{\text{cond}} = 2^k \ln 2 - \frac{3}{2} \ln 2 + o_k(1)$ a much more dramatic type of correlation amongst satisfying assignments than the subtle drift towards σ_{maj} kicks in.

To explain this, we sketch the physics predictions [26] as to the geometry of the set $\mathcal{S}(\Phi)$ of satisfying assignments. According to the cavity method, already for densities $r > 2^k \ln(k)/k$, way below the k -SAT threshold, the set $\mathcal{S}(\Phi)$ has a decomposition

$$\mathcal{S}(\Phi) = \bigcup_{i=1}^{\Sigma} \mathcal{C}_i$$

into an exponential number $\Sigma = \exp(\Omega(n))$ of “clusters” $\mathcal{C}_i$. These clusters are well-separated. That is, any two assignments in different clusters have Hamming distance $\Omega(n)$. Furthermore, within each cluster $\mathcal{C}_i$ most variables (say, at least $0.99n$) are *frozen*, i.e., they take the same truth value under all the assignments in $\mathcal{C}_i$. Finally, each cluster is expected to be internally “well-connected”. That is, one can walk within the cluster $\mathcal{C}_i$ from any $\sigma \in \mathcal{C}_i$ to any other $\tau \in \mathcal{C}_i$ by only altering, say, $\ln n$ variables at each step. The existence of clusters and frozen variables has by now been established rigorously [1, 4, 33].

For $r < r_{\text{cond}}$ each cluster $\mathcal{C}_i$ only contains an $\exp(-\Omega(n))$ fraction of the entire set $\mathcal{S}(\Phi)$ of satisfying assignments. In effect, if we draw two satisfying assignments σ, τ of Φ independently at random, then most likely they belong to different clusters. Thus, we expect σ, τ to have a large Hamming distance. In particular, it is conceivable that they “look uncorrelated”, apart from the inevitable drift towards σ_{maj} . The contribution of [11] was, in a sense, to establish this hunch rigorously.

By contrast, according to the 1RSB cavity method [26], for $r_{\text{cond}} < r < r_{k\text{-SAT}}$ the largest cluster contains a *constant*, i.e., $\Omega(1)$ fraction of the set $\mathcal{S}(\Phi)$. In fact, a bounded number of clusters contain a $1 - o(1)$ fraction of $\mathcal{S}(\Phi)$, a phenomenon called *condensation* in physics jargon. Consequently, if we draw two satisfying assignments σ, τ independently at random, then there is a good chance that σ, τ belong to the same cluster. In that case, they will be *heavily* correlated, because they coincide on all the variable that are frozen in that cluster. Though there is currently no rigorous proof that condensation occurs in random k -SAT, there have been partial rigorous verifications of it in other, symmetric problems [12, 13].

The correlations that condensation induces (or, strictly speaking, is conjectured to induce) not only derail the second moment method, but also the physicists’ “replica symmetric ansatz”. The 1RSB cavity method surmounts this obstacle by switching to a different random variable, namely the number Σ of *clusters*. Provably, Σ must remain exponentially large w.h.p. right up to $r_{k\text{-SAT}}$ [4]. Hence, as clusters are well-separated, there might be a chance that two random clusters decorrelate, even though two randomly chosen satisfying assignments do not. We are going to turn this intuition into a rigorous proof.

To this end, we represent each cluster $\mathcal{C}_i$ by a map $\zeta_i : V \rightarrow \{0, 1, *\}$ in which each variable either takes a Boolean value 0, 1 or the “joker value” $*$. The idea is that $\zeta_i(x_j) = 1$ means that x_j is frozen to the value 1 in the cluster $\mathcal{C}_i$. Similarly, $\zeta_i(x_j) = 0$ indicates that x_j is frozen to 0. By contrast, $\zeta_i(x_j) = *$ means that x_j is unfrozen in the cluster $\mathcal{C}_i$. Fortunately, there is a neat description of the resulting “relaxed assignments” that does not depend on a precise definition of clusters, freezing etc.

DEFINITION 1 ([28]). A map $\zeta : V \rightarrow \{0, 1, *\}$ is a *cover* of a k -CNF $\Phi = \Phi_1 \wedge \dots \wedge \Phi_m$ if the following two conditions are satisfied. Extend ζ to a map from the set of literals to $\{0, 1, *\}$ by letting $\zeta(\neg x_j) = \neg \zeta(x_j)$, with the convention that $\neg 0 = 1$, $\neg 1 = 0$, $\neg * = *$. Then

CV1 each clause Φ_i either contains a literal that takes the value 1 under ζ , or two literals that take the value $*$, and

CV2 any literal l such that $\zeta(l) = 1$ occurs in a clause Φ_i whose other literals are all set to 0.

In terms of the cluster intuition, condition **CV1** provides that each clause either contains one literal that is frozen to true, or at least two that are unfrozen. (The idea is that no unfrozen literal l may occur in a clause whose other $k-1$ literals are frozen to 0, as otherwise the clause would freeze l to 1.) In addition, **CV2** ensures that each variable x_j that takes the value 0 or 1 is frozen to this value because there is a clause Φ_i whose other $k-1$ literals are frozen to values that do not satisfy Φ_i . Hence, we expect that the clusters and covers of Φ are (essentially) in one-to-one correspondence, and our proof vindicates this notion.

Thus, the programme is to perform a second moment argument for the number of covers.⁴ Yet matters are far from straightforward as the asymmetry of the k -SAT problem implies that, much like satisfying assignments, covers lean towards σ_{maj} and thus are subtly correlated. In effect, a “vanilla” second moment argument cannot succeed.

To accommodate the drift towards σ_{maj} , we employ the physicists’ Survey Propagation technique. Survey Propagation is a message passing procedure for (heuristically) calculating the marginal probability that a fixed variable x_j takes each value 0, 1, $*$ in a random cover ζ of Φ [5, 31]. The details of Survey Propagation are rather non-trivial (e.g., they involve solving a seriously complicated fixed point problem on the space of probability measures on the 3-simplex), and the result is not explicit. However, asymptotically the dominant terms result from the degrees $d_{x_j}, d_{\neg x_j}$. Indeed, for densities $r_{\text{cond}} < r < r_{k\text{-SAT}}$ Survey Propagation predicts that

$$\begin{aligned} \mathbb{P}[\zeta(x_j) = z \mid d_{x_j}, d_{\neg x_j}] &= \\ &\begin{cases} \frac{1}{2} + \frac{d_{x_j} - d_{\neg x_j}}{2^{k+1}} - 2^{-k-2} + o_k(2^{-k}) & \text{if } z = 1, \\ \frac{1}{2} - \frac{d_{x_j} - d_{\neg x_j}}{2^{k+1}} - 2^{-k-2} + o_k(2^{-k}) & \text{if } z = 0, \\ \frac{1}{2} - \frac{d_{x_j} - d_{\neg x_j}}{2^{k+1}} - 2^{-k-2} + o_k(2^{-k}) & \text{if } z = *. \end{cases} \end{aligned} \quad (5)$$

The probability term on the left hand side refers to choosing a random formula Φ and then a random cover ζ of Φ , given the degrees of two literals $x_j, \neg x_j$. The approximation (5) is valid so long as $|d_{x_j} - d_{\neg x_j}| = o_k(2^k)$, a condition that holds for the vast majority of variables w.h.p. Observe that (5) is very much in line with our intuition that covers lean towards σ_{maj} . In Section 4 we are going to craft a random variable around (5) that allows us to incorporate this drift, and thus to make the second moment argument work.

⁴Dimitris Achlioptas suggested the general strategy of applying the second moment method to “covers” as early as 2007/8. But at the time it was not clear what the appropriate definition of covers might be, nor how to carry out such a second moment argument.

2.4 Random regular k -SAT

In the case of the uniformly random formula Φ , it seems infeasible to solve the (distributional) fixed point equations resulting from Survey Propagation explicitly. However, in a simpler model of random formulas, a more explicit calculation is possible. Namely, let $\Phi_{k,d\text{-reg}}$ denote a k -CNF on the variables $V = \{x_1, \dots, x_n\}$ in which each of the $2n$ literals $x_1, \neg x_1, \dots, x_n, \neg x_n$ occurs exactly d times, chosen uniformly at random among all such formulas. Hence, $d_{x_i} = d_{\neg x_i} = d$ for all i . In this model, there clearly is no drift towards the (trivial) majority vote assignment. This implies that the Survey Propagation calculation boils down to solving a certain fixed point problem on the 3-simplex (rather than the set of probability measures on the 3-simplex).

This leads to the following result. Observe that *every* formula has got a cover, namely the trivial one that sets all variables to $*$. Of course, this cover is utterly unhelpful in our quest for satisfying assignments. In fact, the cavity method predicts that near the k -SAT threshold all clusters correspond to covers with no more than $2^{-k}n$ variables set to $*$. Thus, let $\Sigma'(\Phi_{k,d\text{-reg}})$ be the number of covers of the random formula $\Phi_{k,d\text{-reg}}$ with at most $2^{-k}n$ variables assigned $*$, and let

$$\Xi(k, d) = \lim_{n \rightarrow \infty} \frac{1}{n} \ln \mathbb{E}[\Sigma'(\Phi_{k,d\text{-reg}})].$$

This limit provably exists, and Survey Propagation predicts that $\Phi_{k,d\text{-reg}}$ is satisfiable w.h.p. if k, d are such that $\Xi(k, d) > 0$, and unsatisfiable w.h.p. if $\Xi(k, d) < 0$.

THEOREM 2. *There is a constant $k_0 \geq 3$ such that the following is true for all $k \geq k_0$.*

1. *If d is such that $\Xi(k, d) \geq 0$, then $\Phi_{k,d\text{-reg}}$ has an assignment $\sigma : V \rightarrow \{0, 1\}$ that satisfies all but $o(n)$ clauses w.h.p.*
2. *If d is such that $\Xi(k, d) < 0$, then w.h.p. under any assignment $\sigma : V \rightarrow \{0, 1\}$ at least $\Omega(n)$ clauses are unsatisfied.*

REMARK 1. *In the first part of Theorem 2, we obtain an assignment that satisfies a $1 - o(1)$ -fraction of all clauses rather than an actual satisfying assignment. This is because there is no counterpart to Lemma 1 in random regular formulas. The proof of Theorem 2 actually shows that if $\Xi(k, d) > 0$, then $\liminf_{n \rightarrow \infty} \mathbb{P}[\Phi_{k,d\text{-reg}} \text{ is satisfiable}] > 0$.*

2.5 Summary and perspective

In summary, we establish Theorem 1 via a new “1RSB-tight” second moment argument, based on counting covers that are tilted towards the majority assignment as suggested by Survey Propagation. By comparison to the replica symmetric approach from [11] (based on counting satisfying assignments), here we face several new challenges. They derive from the fact that the notion of “cover” is inherently more sophisticated than the simple concept of a satisfying assignment. For instance, in covers there are three possible “spins” 0, 1, $*$, out of which one ($*$) occurs far less frequently than the others. This substantially increases the complexity of the formulas that we need to deal with. But more importantly, while condition **CV1** is broadly similar to the notion of “satisfying assignment”, condition **CV2** imposes a completely new kind of constraint. Indeed, **CV2** introduces an

occupancy problem into our analysis, which means that critical clauses play a very special role. To deal with this, we are going to introduce a color code (see Section 4 below), which further increases the number of variables.

Theorem 1 determines the k -SAT threshold up to an error that vanishes as k gets larger. While no attempt has been made to optimize the error term, it is easily bounded by $\exp(-\Omega_k(k))$. Thus, Theorem 1 clearly leaves open the problem of calculating the exact threshold for any given k (say, $k = 3$). But it seems that one cannot expect a fully explicit answer to this problem. As mentioned, the physics prediction with respect to this problem is in terms of the solution to a fixed point problem on the set of probability measures on the 3-simplex, and this fixed point problem may well be computationally intractable.

A further interesting direction might be an analysis of *Survey Propagation Guided Decimation*, the physicists' flagship algorithm for solving random k -SAT instances. This algorithm combines the Survey Propagation message passing procedure with a decimation step, where one variable is selected, set to a specific truth value and eliminated from the formula. Of course, just because Survey Propagation predicts the correct k -SAT threshold doesn't mean that the algorithm succeeds up to $r_{k\text{-SAT}}$. Indeed, the algorithms' success seems to hinge on a rapid decay of correlations between covers (in a certain well-defined sense) even as the algorithm eliminates one variable after the other. In fact, it seems questionable that this property remains valid throughout the decimation process for densities near the k -SAT threshold. Thus, it would be interesting to see if the present techniques can be used to figure out for what densities r this correlation decay occurs.

3. RELATED WORK

This is one of the first papers to (at least asymptotically) vindicate the 1RSB cavity method rigorously, and the first to do so in an asymmetric problem. In [10] we obtained a result similar to Theorem 1 for the (symmetric) random k -NAESAT problem. In symmetric problems many of the maneuvers that we are going to have to go through (e.g., clause/variable types, see Section 4) are unnecessary. Independently, Ding, Sly and Sun [14, 15] verified the 1RSB prediction exactly in the random regular k -NAESAT problem (where each variable appears exactly d times), and in the independent set problem in random regular graphs. Both of these problems are symmetric. The proofs in [14, 15] are based on the second moment method applied to a notion of "cover" appropriate for NAESAT/independent sets, while [10] relies on an ad-hoc concept called "heavy solutions".

All other random CSPs where the threshold for the existence of solutions is known exactly are replica symmetric (in the sense that condensation does not occur). For instance, in the random k -SAT problem with $k > \log_2 n$ (i.e., the clause length tends to infinity with the number of variables), a "vanilla" second moment argument suffices to identify the threshold [20]. Another example is the random k -XORSAT problem (random linear equations mod 2) [17, 35]. Furthermore, the exact satisfiability threshold is known in random 2-SAT [7, 21]. This is, of course, a special case, as 2-SAT admits a simple criterion for (un)satisfiability, on which the proofs of [7, 21] hinge. In several other examples the replica

symmetric predictions have been validated rigorously (see, e.g., [31, Chapter 15–17]).

As mentioned earlier, the best prior bounds on the k -SAT threshold were obtained by far simpler second moment arguments. Achlioptas and Moore [2] got within (about) a factor of two of the k -SAT threshold. They actually apply the second moment method to the number of NAE-satisfying assignments, thereby symmetrizing the problem. Subsequently, Achlioptas and Peres [3] discovered a more subtle way to guarantee symmetry ("balanced assignments"), which allowed them to get within an additive $\frac{k}{2} \ln 2 + O_k(1)$ of $r_{k\text{-SAT}}$. In addition, in our prior work [11] we verified the replica symmetric lower bound (1) on $r_{k\text{-SAT}}$ by designing a second moment argument that can accommodate asymmetry. The present work builds upon some of the ideas from [11] such as clause/literal types, but substantial new ingredients are necessary to cope with the 1RSB scenario. The random regular k -SAT problem was previously studied via the "vanilla" second moment method by Rathi, Aurell, Rasmussen and Skoglund [36]. They got close to the replica symmetric lower bound in that problem, while Theorem 2 verifies the 1RSB prediction. As mentioned above, Theorem 1 does not improve the current lower bound on $r_{3\text{-SAT}}$, which remains 3.52 [22].

The best current algorithms for random k -SAT find satisfying assignments w.h.p. for densities up to $\approx 1.817 \cdot 2^k/k$ (better for small k) resp. $2^k \ln k/k$ (better for large k) [8, 19], a factor of $\Theta(k/\ln k)$ below the satisfiability threshold.

The notion of covers, which plays a key role in the 1RSB cavity method, has so far received only limited attention in rigorous work. In an important conceptual contribution, Maneva, Mossel and Wainwright [27] introduced a similar concept ("core assignments") to show that (generalized) Survey Propagation can be viewed as Belief Propagation on a modified Markov random field. Furthermore, Maneva and Sinclair [28] used covers to prove a (conditional) upper bound on the 3-SAT threshold in uniformly random formulas.

4. THE RANDOM VARIABLE

The aim in this section is to design the random variable upon which the proof of Theorem 1 is based. Throughout, we assume that $r = 2^k \ln 2 - \frac{1+\ln 2}{2} - \varepsilon_k$ with ε_k tending to 0 slowly in the limit of large k .

4.1 Pruning

As a first step, we are going to rid the formula of all the literals whose degrees deviate substantially from the average degree $\frac{km}{2n} \sim \frac{1}{2}kr$. To accomplish this, we prune the formula as follows.

PR1 Initially, let U be the set of all variables $x \in V$ such that

$$\max\{|d_x - kr/2|, |d_{\neg x} - kr/2|\} > k^3 2^{k/2}. \quad (6)$$

PR2 While there is a clause that features at least three variables from U ,

- remove all such clauses from the formula, and
- add all variables x to U that satisfy (6) in the reduced formula.

PR3 Finally, remove the variables in U from all the remaining clauses.

Let Φ' denote the formula obtained from Φ via **PR1**–**PR3** and let $V' = V \setminus U$ be its variable set. We may assume without loss of generality that $V' = \{x_1, \dots, x_{n'}\}$. Moreover, let $d'_x, d'_{\neg x}$ be the degrees of the literals $x, \neg x$ for $x \in V'$ in Φ' . Standard arguments yield

PROPOSITION 1. *W.h.p. the random formula Φ has the following properties.*

1. *If σ' is a satisfying assignment of Φ' , then Φ has a satisfying assignment σ such that $\sigma(x) = \sigma'(x)$ for all $x \in V'$.*
2. *We have $\sum_{x \notin V'} d_x + d_{\neg x} \leq \exp(-k^2)n$.*

Thus, it suffices to show that Φ' is satisfiable w.h.p. By construction, we have

$$|d'_x - \frac{1}{2}kr|, |d'_{\neg x} - \frac{1}{2}kr| \leq k^3 2^{k/2} \quad \text{for all } x \in V'. \quad (7)$$

Moreover, all the clauses $\Phi'_1, \dots, \Phi'_{m'}$ of Φ' have length either $k-2$ or $k-1$ or k . Indeed, the second part of Proposition 1 implies that the vast majority have length k . By the principle of deferred decisions, we can characterize the distribution of the random formula Φ' as follows.

FACT 1. *Conditional on n', m' , the degree sequence $d' = (d'_x, d'_{\neg x})_{x \in V'}$, and the clause lengths $k' = (k'_1, \dots, k'_{m'})$, the formula Φ' is uniformly random.*

In the light of Fact 1, we can think of generating the random formula Φ' as follows. The process **PR1**–**PR3** induces a probability distribution D' over (d', k') , i.e., the degree sequence and the clause lengths of the pruned formula. For a given pair (d', k') , let $\Phi^{d', k'}$ be a formula with degree sequence d' and clause lengths k' chosen uniformly at random. Then we can mimic the distribution of the pruned formula Φ' by first choosing (d', k') from the distribution D' , and then generating $\Phi^{d', k'}$. Therefore,

$$\mathbb{P}[\Phi' \text{ is satisfiable}] = \mathbb{E}_{(d', k') \in D'} [\mathbb{P}[\Phi^{d', k'} \text{ is satisfiable}]].$$

Hence, by Lemma 1 and Proposition 1, it suffices to prove that

$$\liminf_{n \rightarrow \infty} \mathbb{E}_{(d', k') \in D'} [\mathbb{P}[\Phi^{d', k'} \text{ is satisfiable}]] > 0. \quad (8)$$

4.2 The color code

As laid out in Section 2, we are going to establish (8) by performing a second moment argument over the number of covers. By comparison to satisfying assignments, covers involve one significant intricacy that we need to deal with. Namely, while condition **CV1** in the definition of “cover” is similar in spirit to the notion of “satisfying assignment”, **CV2** imposes the additional requirement that each literal set to 1 be “frozen”. In effect, **critical clauses**, i.e., clauses that contain one literal set to 1 while all other literals are set to 0, play a special role: each literal set to true must occur in one of them.

To implement this, we need to get a better grip on $\Phi^{d', k'}$. We employ a construction reminiscent of the “configuration model” of random graphs. With

$$L' = \{x_1, \neg x_1, \dots, x_{n'}, \neg x_{n'}\}$$

the set of literals, let

$$\mathcal{L}^{d'} = \bigcup_{l \in L'} \{l\} \times [d'_l]$$

be a set that contains d'_l “clones” of each literal l . Moreover, let

$$\mathcal{I}^{d', k'} = \bigcup_{i=1}^{m'} \{i\} \times [k'_i]$$

be the set of all “literal slots” in the formula $\Phi^{d', k'}$. Then we can think of $\Phi^{d', k'}$ simply as a random bijection

$$\Phi^{d', k'} : \mathcal{I}^{d', k'} \rightarrow \mathcal{L}^{d'}, \quad (i, j) \mapsto \Phi_{ij}^{d', k'}$$

that specifies which clone occupies which slot.

To deal with the special role that critical clauses play, we introduce a color code. If ζ is a cover of $\Phi^{d', k'}$, then we extend ζ to a map ξ from the set $\mathcal{L}^{d'}$ of clones to the colors red, blue, green, yellow (**r, b, g, y**, for short). The semantics of the colors is this. All clones of literals that are set to $*$ under ζ are colored green, and all clones of literals that ζ assigns 0 are colored yellow. Moreover, clones that are set to 1 under ζ are colored either red or blue: those that occur in critical clauses are red, the others blue. The colorings that emerge in this way admit a neat characterization.

DEFINITION 2. *A map $\xi : \mathcal{L}^{d'} \rightarrow \{\mathbf{r}, \mathbf{b}, \mathbf{g}, \mathbf{y}\}$ is a **shade** if the following conditions are satisfied.*

SD1 *For any literal $l \in L'$ one of the following statements is true.*

- *all clones of l and $\neg l$ are colored green under ξ .*
- *all clones of l are colored either red or blue, and all clones of $\neg l$ are colored yellow.*
- *all clones of l are colored yellow, and all clones of $\neg l$ are colored either red or blue.*

SD2 *There is no literal $l \in L'$ all of whose clones are blue.*

Condition **SD2** is to ensure that a literal set to 1 is “frozen” by a critical clause, represented by a red clone.

In terms of this color code, we can express easily that a shade ξ corresponds to a cover. Indeed, we call a shade ξ **valid** if the following two conditions are satisfied.

V1 *If a clause contains a red clone, then all its other clones are colored yellow.*

V2 *Any clause without a red clone contains at least two clones colored blue or green.*

With this terminology in place, valid shades of $\Phi^{d', k'}$ are in one-to-one correspondence with covers, and the concept of shades allows us to track the critical clauses of $\Phi^{d', k'}$ explicitly.

4.3 The majority vote

We still need to accommodate the drift towards the majority vote assignment. To specify the desired drift, we are going to prescribe explicitly for each literal its “preference” for each of the colors **r, b, g, y**. To this end, we consider a map θ from the integers $\mathbf{Z}$ to the probability distributions over the colors $\{\mathbf{r}, \mathbf{b}, \mathbf{g}, \mathbf{y}\}$. Formally,

$$\theta : \begin{cases} \mathbf{Z} \rightarrow \{(\theta_{\mathbf{r}}, \theta_{\mathbf{b}}, \theta_{\mathbf{g}}, \theta_{\mathbf{y}}) \in [0, 1]^4 : \theta_{\mathbf{r}} + \theta_{\mathbf{b}} + \theta_{\mathbf{g}} + \theta_{\mathbf{y}} = 1\}, \\ z \mapsto \theta(z) = (\theta_{\mathbf{r}}(z), \theta_{\mathbf{b}}(z), \theta_{\mathbf{g}}(z), \theta_{\mathbf{y}}(z)). \end{cases}$$

Together with the degree sequence d' , θ induces a map on the set L' of literals, namely

$$\theta^{d'}(l) = \theta(d'_l - d'_{\neg l}). \quad (9)$$

Thus, the preference that a literal l has for each color is governed by the difference $d_l - d_{-l}$. Of course, this notion is inspired by the asymptotic expansion of the Survey Propagation fixed point (5).

In fact, to come up with a good choice of θ , we employ the Survey Propagation prediction (5). But we need to turn (5) into a formula that incorporates the four colors $\mathbf{r}, \mathbf{b}, \mathbf{g}, \mathbf{y}$. That is, we need to split the probability mass assigned to 1 in (5) into probabilities for red and blue. A simple calculation based on the Survey Propagation formalism suggests that the mass assigned to red should be approximately just the constant 2^{-k} . Hence, we arrive at

$$\theta_c(z) = \begin{cases} 2^{-k} & \text{if } c = \mathbf{r}, \\ \frac{1}{2} + 2^{-k-1}z - 5 \cdot 2^{-k-2} & \text{if } c = \mathbf{b}, \\ \frac{1}{2} - 2^{-k-1}z - 2^{-k-2} & \text{if } c = \mathbf{y}, \\ 2^{-k-1} & \text{if } c = \mathbf{g}, \end{cases} \quad (10)$$

provided that $|z| \leq k^3 2^{1+k/2}$. Due to (7), the argument $z = d_l - d_{-l}$ in (9) always satisfies $|z| \leq k^3 2^{1+k/2}$.

4.4 Literal and clause types

Let us call $\theta^{d'}(l)$ the *type of the literal* l , and let

$$\mathcal{T}^{d'} = \{\theta^{d'}(l) : l \in L'\}$$

be the set of all possible types. Additionally, define the *type of a clause* $l_1 \vee \dots \vee l_h$ as the tuple

$$(\theta^{d'}(l_1), \dots, \theta^{d'}(l_h))$$

comprising the types of the literals. Let $\mathcal{T}_*^{d'}$ denote the set of all possible clause types. For each clause $\Phi_i^{d',k'}$ let λ_i denote its type, and let $\lambda = (\lambda_1, \dots, \lambda_{m'})$ comprise the types of all clauses.

To proceed, we need to condition on the clause types. To this end, we use a similar trick as above when we fixed the degree sequence. That is, we think of creating the random formula $\Phi^{d',k'}$ in two steps. First, choose a vector $\lambda = (\lambda_1, \dots, \lambda_{m'})$ of clause types from the distribution λ . Then, generate a formula $\Phi^{d',\lambda}$ among all formulas with literal degrees d' and clause types as detailed by λ uniformly at random. By construction,

$$\mathbb{P}[\Phi^{d',k'} \text{ is sat}] = \mathbb{E}_{\lambda \in \lambda} [\mathbb{P}[\Phi^{d',\lambda} \text{ is sat}]]. \quad (11)$$

Thus, to establish (8) we are going to show that

$$\liminf_{n \rightarrow \infty} \mathbb{E}_{d',\lambda} [\mathbb{P}[\Phi^{d',\lambda} \text{ is satisfiable}]] > 0.$$

The random formula $\Phi^{d',\lambda}$ has a very simple combinatorial description. Given the vector λ , we know exactly which clause requires how many literals of what type. Thus, to generate the formula, we randomly match each literal slot in the formula to a clone of the required type.

With the types of the literals and clauses in place, we can formalize what it means for a shade to respect the preferences expressed by θ . Let $L'(d^+, d^-)$ be the set of all literals $l \in L'$ with $d'_l = d^+$ and $d'_{-l} = d^-$. Then we say that a shade ξ is a *θ -shade* if for all d^+, d^- and all colors $c \in \{\mathbf{r}, \mathbf{b}, \mathbf{g}, \mathbf{y}\}$ out of all the (typically $\Omega(n)$) clones of the literals in $L'(d^+, d^-)$, precisely a $\theta_c(d^+ - d^-)$ fraction have

color c . In symbols,

$$\sum_{l \in L'(d^+, d^-)} \sum_{h=1}^{d^+} \mathbf{1}_{\xi(l,h)=c} \propto \theta_c(d^+ - d^-),$$

where the $\propto$ sign hides the normalizing factor $d^+ |L'(d^+, d^-)|$.

Similarly, we call ξ *judicious* if for each clause type $\ell = (t^1, \dots, t^h)$, each $j \in [h]$ and every color $c \in \{\mathbf{r}, \mathbf{b}, \mathbf{g}, \mathbf{y}\}$, out of all the clauses of type ℓ , exactly a t_c^j fraction have a clone colored c in their j th position. In symbols,

$$\sum_{i=1}^{m'} \mathbf{1}_{\lambda_i = \ell, \xi(\Phi_{ij}^{d',\lambda}) = c} \propto t_c^j,$$

where the normalizing term is the number of clauses of type ℓ . The random variable that the proof is based on is going to be essentially the number Z'' of valid, judicious θ -shades. With a bit of effort, it is possible to compute its first moment explicitly. This yields

PROPOSITION 2. *With high probability over the choice of d', λ we have $\mathbb{E}[Z''(\Phi^{d',\lambda})] = \exp(\Omega(n))$.*

4.5 Back to satisfying assignments

The random variable Z'' counts valid shades which, as we saw, correspond to covers. Indeed, by construction we can obtain a cover $\hat{\xi}$ from a valid shade ξ by letting $\hat{\xi}(x) = 1$ if $\xi(x, 1) \in \{\mathbf{r}, \mathbf{b}\}$, $\hat{\xi}(x) = 0$ if $\xi(x, 1) = \mathbf{y}$, and $\hat{\xi}(x) = *$ if $\xi(x) = \mathbf{g}$. But we still need to make sure that this cover can actually be extended to a satisfying assignment. Thus, let us call ξ *extendible* if it is possible to set the variables assigned $*$ under $\hat{\xi}$ to either 0 or 1 so as to obtain a satisfying assignment. Let Z' be the number of extendible, valid, judicious θ -shades.

PROPOSITION 3. *W.h.p. over the choice of d', λ we have*

$$\mathbb{E}[Z'(\Phi^{d',\lambda})] = (1 + o(1)) \mathbb{E}[Z''(\Phi^{d',\lambda})].$$

PROOF. We condition on the event that ξ is a valid, judicious θ -shade. To extend $\hat{\xi}$ to a satisfying assignment, we only need to worry about those clauses that do not contain a literal l with $\hat{\xi}(l) = 1$. By condition **V2**, each such clause contains at least two clones colored green. Thus, we can obtain a 2-SAT formula by selecting two green clones in each such clause randomly. A simple calculation shows that typically the number of such 2-clauses is $O_k(4^{-k})n$, while the number of variables that $\hat{\xi}$ sets to $*$ is $\Theta_k(2^{-k})n$. Thus, the 2-SAT instance is extremely sparse, and arguments along the lines of [7, 21] show that it is satisfiable w.h.p. $\square$

4.6 Separability

We impose one last condition to facilitate the second moment argument. Namely, according to the physics picture, each cluster corresponds to a single cover, and any two covers are well-separated. To hard-wire this geometry into our random variable, we call a valid θ -shade ξ *separable* if the cover $\hat{\xi}$ that it gives rise to has the following property:

The total number of covers ζ such that the Hamming distance between ζ and $\hat{\xi}$ does not lie in the interval $[\frac{n}{2}(1 - k^4 2^{-k/2}), \frac{n}{2}(1 + k^4 2^{-k/2})]$ is $\exp(o(n))$.

Finally, let $\mathcal{Z}$ be the number of separable, extendible, valid, judicious, or, for the sake of brevity, *good* θ -shades. The expansion properties of the random formula yield

PROPOSITION 4. *W.h.p. d', λ are such that*

$$\mathbb{E}[\mathcal{Z}(\Phi^{d', \lambda})] = (1 + o(1))\mathbb{E}[\mathcal{Z}''(\Phi^{d', \lambda})] = \exp(\Omega(n)).$$

To prove Theorem 1, we apply the second moment method to $\mathcal{Z}$.

5. THE SECOND MOMENT

The second moment $\mathbb{E}[\mathcal{Z}(\Phi^{d', \lambda})^2]$ is nothing but the expected number of *pairs* of good θ -shades. Therefore, with ξ_1, ξ_2 ranging over all θ -shades, establishing the second moment bound (3) is equivalent to showing

$$\sum_{\xi_1, \xi_2} \mathbb{P}[\xi_1, \xi_2 \text{ are good in } \Phi^{d', \lambda}] \leq C \cdot \mathbb{E}[\mathcal{Z}(\Phi^{d', \lambda})]^2 \quad (12)$$

for some n -independent number $C > 0$. As mentioned earlier, this really amounts to establishing is that the sum on the l.h.s. is dominated by “uncorrelated” ξ_1, ξ_2 . Indeed, the purpose of the concept of types in the previous section was to define a “slice” inside the set of all possible shades so that inside this slice, “uncorrelated” means that both ξ_1, ξ_2 have the drift towards the majority vote prescribed by θ . This slice is the set of judicious θ -shades.

To turn this intuition into a proof, we need a measure of how “correlated” two θ -shades ξ_1, ξ_2 are. To this end, we define for each type $t \in \mathcal{T}^{d'}$ and any two colors $a, b \in \{\mathbf{r}, \mathbf{b}, \mathbf{g}, \mathbf{y}\}$

$$\mathcal{O}_t^{ab}(\xi_1, \xi_2) \propto \sum_{l \in L'} \sum_{j=1}^{d_l} \mathbf{1}_{\theta^{d'}(l)=t, \xi_1(l,j)=a, \xi_2(l,j)=b},$$

where the normalization is such that for

$$\sum_{a,b \in \{\mathbf{r}, \mathbf{b}, \mathbf{g}, \mathbf{y}\}} \mathcal{O}_t^{ab} = 1 \quad \text{for each } t.$$

In words, $\mathcal{O}_t^{ab}(\xi_1, \xi_2)$ is the fraction of clones of type t that have color a under ξ_1 and color b under ξ_2 . We call

$$\mathcal{O}(\xi_1, \xi_2) = (\mathcal{O}_t^{ab}(\xi_1, \xi_2))_{t,a,b}$$

the **overlap** of ξ_1, ξ_2 . Now, what overlap corresponds to “uncorrelated” ξ_1, ξ_2 ? A straightforward calculation yields

LEMMA 2. *Choose two θ -shades ξ_1, ξ_2 uniformly and independently. Then $\mathbb{E}[\mathcal{O}_t^{ab}(\xi_1, \xi_2)] = t_a \cdot t_b$.*

Thus, the “uncorrelated” overlap of two θ -shades ξ_1, ξ_2 is $\bar{\mathcal{O}} = (\bar{\mathcal{O}}_t^{a,b})_{t,a,b}$ with $\bar{\mathcal{O}}_t^{a,b} = t_a \cdot t_b$; we emphasize that in Lemma 2 ξ_1, ξ_2 are *not* required to be valid. The core of the second moment argument consists in establishing

PROPOSITION 5. *There is $C = C(k) > 0$ such that with high probability over the choice of d', λ we have*

$$\mathbb{E}[\mathcal{Z}(\Phi^{d', \lambda})^2] \leq C \cdot \sum_{\xi_1, \xi_2: \|\mathcal{O}(\xi_1, \xi_2) - \bar{\mathcal{O}}\|_\infty \leq n^{-1/2}} \mathbb{P}[\xi_1, \xi_2 \text{ are good}].$$

For ξ_1, ξ_2 such that

$$\|\mathcal{O}(\xi_1, \xi_2) - \bar{\mathcal{O}}\|_\infty \leq n^{-1/2},$$

i.e., “uncorrelated” ξ_1, ξ_2 , it is not difficult to calculate the probability that both ξ_1, ξ_2 are good. In fact, it is almost immediate from the construction that this is, up to a constant factor, just the product of the probabilities that ξ_1, ξ_2 are good individually. Thus, (12) is immediate from Proposition 5.

To prove Proposition 5, we consider the function

$$\mathcal{F}(\mathcal{O}) = \frac{1}{n} \ln \sum_{\xi_1, \xi_2: \mathcal{O}(\xi_1, \xi_2) = \mathcal{O}} \mathbb{P}[\xi_1, \xi_2 \text{ are good}] \quad (13)$$

that captures the contribution of a given overlap $\mathcal{O}$ to (12). We work on a logarithmic scale here because both the number of θ -shades ξ_1, ξ_2 with $\mathcal{O}(\xi_1, \xi_2) = \mathcal{O}$ and the probability term scale exponentially in n . In fact, a moment of reflection reveals that the probability term depends only on the overlap $\mathcal{O}$ but not on the actual shades ξ_1, ξ_2 . Thus, letting

$$\begin{aligned} \mathcal{N}(\mathcal{O}) &= \frac{1}{n} \ln |\{(\xi_1, \xi_2) : \mathcal{O}(\xi_1, \xi_2) = \mathcal{O}\}|, \\ \mathcal{P}(\mathcal{O}) &= \frac{1}{n} \ln \mathbb{P}[\text{both } \xi_1, \xi_2 \text{ are good in } \Phi^{d', \lambda}] \end{aligned}$$

for any ξ_1, ξ_2 with $\mathcal{O}(\xi_1, \xi_2) = \mathcal{O}$, we can rewrite (13) as

$$\mathcal{F}(\mathcal{O}) = \mathcal{N}(\mathcal{O}) + \mathcal{P}(\mathcal{O}). \quad (14)$$

By standard arguments (“Laplace method”), to prove Proposition 5 it suffices to show that the function $\mathcal{F}(\mathcal{O})$ attains its global maximum at $\bar{\mathcal{O}}$.

This is easily seen to be true of the “entropy term” $\mathcal{N}(\mathcal{O})$. In fact, together with standard large deviations arguments, Lemma 2 shows that $\mathcal{N}(\mathcal{O})$ is a concave function that takes its global maximum at $\bar{\mathcal{O}}$. By contrast, the function $\mathcal{P}(\mathcal{O})$ takes its global maximum at the point $\mathcal{O}_t^{ab} = \mathbf{1}_{a=b} t_a$, which simply corresponds to $\xi_1 = \xi_2$; for the probability that both ξ_1, ξ_2 are good is maximized if $\xi_1 = \xi_2$.

Thus, to prove that $\mathcal{F}(\mathcal{O})$ takes its global maximum at $\bar{\mathcal{O}}$, we need to

SMM1 show that $\bar{\mathcal{O}}$ is a stationary point⁵ of the probability term $\mathcal{P}(\mathcal{O})$,

SMM2 expand both $\mathcal{N}(\mathcal{O})$ and $\mathcal{P}(\mathcal{O})$ around $\bar{\mathcal{O}}$ to verify that $\bar{\mathcal{O}}$ is a local maximum,

SMM3 verify that there are no other local maxima with a function value greater than $\mathcal{F}(\bar{\mathcal{O}})$.

Out of these three tasks, **SMM1** is the crucial one conceptually, while the others are of a (difficult) technical nature. Indeed, **SMM1** is a necessary condition for our entire construction (shades, literal types, clause types, ...) to yield a random variable under which pairs of covers decorrelate.

Thus, we conclude with a discussion of **SMM1**. For starters, we need to write $\mathcal{P}(\mathcal{O})$ in a more explicit form. Unfortunately, $\mathcal{O}$ is still too coarse a parameter to enable this. To define a more refined overlap parameter, let ξ_1, ξ_2 be two θ -shades with $\mathcal{O}(\xi_1, \xi_2) = \mathcal{O}$. As we saw above, we can think of the formula $\Phi^{d', \lambda}$ as a type-preserving random matching between the sets $\mathcal{L}^{d'}$ of clones the set $\mathcal{T}^{d', k'}$ of “slots” in the formula. We need to keep track of how this random matching distributes the clones of various colors to

⁵Strictly speaking, $\mathcal{P}(\mathcal{O})$ is a function defined on a discrete domain, but it can be extended to a continuous domain canonically.

the clauses. To this end, for any clause type $\ell = (t^1, \dots, t^h)$, any $j \in [h]$ and any two colors a, b we let $\omega_{\ell,j}^{a,b}(\xi_1, \xi_2)$ be the fraction of clauses of $\Phi^{d',\lambda}$ of type ℓ whose j th literal has color a under ξ_1 and color b under ξ_2 . Moreover, we set

$$\omega(\xi_1, \xi_2) = (\omega_{\ell,j}^{a,b}(\xi_1, \xi_2))_{\ell,j,a,b}.$$

Thus, $\omega(\xi_1, \xi_2)$ comprises the “joint color statistics” broken down to the individual clause types. Let $\Omega(\mathcal{O})$ be the set of all possible outcomes of the random variable $\omega(\xi_1, \xi_2)$ with $\mathcal{O}(\xi_1, \xi_2) = \mathcal{O}$. For a given $\omega \in \Omega(\mathcal{O})$, we let

$$\mathcal{P}(\omega) = \frac{1}{n} \ln \mathbb{P}[\xi_1, \xi_2 \text{ are good} | \omega(\xi_1, \xi_2) = \omega].$$

LEMMA 3. *Let $\bar{\omega}$ be the vector with entries $\bar{\omega}_{\ell,j}^{c,c'} = \bar{\mathcal{O}}_{\ell,j}^{c,c'}$ for any clause type $\ell = (t^1, \dots, t^h)$. Then $\bar{\omega}$ is a stationary point of $\mathcal{P}(\omega)$.*

PROOF. We prove this statement by a combinatorial argument, i.e., without actually deriving explicit formulas for $\mathcal{P}(\omega)$ and differentiating them. The proof, which is an extension of an argument from [11] to the 1RSB scenario, illustrates why it is important to work with judicious shades. Let ξ_1, ξ_2 denote two θ -shades that are chosen independently and uniformly at random.

Let m_ℓ be the number of clauses of type ℓ . The random formula $\Phi^{d',\lambda}$ provides a matching between the literal slots in the formula and the clones, which are colored by ξ_1, ξ_2 . Thus, for any clause $i \in [m_\ell]$ and any j let $\psi_1(i, j)$ denote the color of the j th literal in the i th clause under ξ_1 . Similarly, let $\psi_2(i, j)$ denote the color of that clone under ξ_2 . Of course, both $\psi_1(i, j)$, $\psi_2(i, j)$ are random variables, determined by ξ_1, ξ_2 and the random matching $\Phi^{d',\lambda}$.

Let $\Psi(\omega)$ be the set of all possible pairs (ψ_1, ψ_2) that can result from this experiment under the condition that $\omega(\xi_1, \xi_2) = \omega$. Thus, $\Psi(\omega)$ contains all pairs (ψ_1, ψ_2) with the following property.

For any clause type $\ell = (t^1, \dots, t^h)$ and any $j \in [h]$ the fraction of clauses of type ℓ such that $\psi_1(i, j) = c$ and $\psi_2(i, j) = c'$ is precisely $\omega_{\ell,j}^{c,c'}$.

If we condition on $\omega(\xi_1, \xi_2) = \omega$, then $(\psi_1, \psi_2) \in \Psi(\omega)$ is a uniformly random element.

Therefore, we can cast $\mathcal{P}(\omega)$ as follows. Let $G(\omega)$ be the set of all outcomes $(\psi_1, \psi_2) \in \Psi(\omega)$ that are possible if ξ_1, ξ_2 are good. Then

$$\mathcal{P}(\omega) = \frac{1}{n} \ln \frac{|G(\omega)|}{|\Psi(\omega)|} = \frac{1}{n} \ln |G(\omega)| - \frac{1}{n} \ln |\Psi(\omega)|. \quad (15)$$

To show that $\bar{\omega}$ is a stationary point of $\mathcal{P}(\omega)$, we are going to argue that $\omega = \bar{\omega}$ is the maximizer of both $|G(\omega)|$ and $|\Psi(\omega)|$. Indeed, by the very definition of $\bar{\omega}$, we have

$$\mathbb{E}[\omega(\xi_1, \xi_2)] = \bar{\omega}.$$

Furthermore, by standard arguments (e.g., Azuma’s inequality), $\omega(\xi_1, \xi_2)$ is tightly concentrated about $\bar{\omega}$. As a result, $\omega = \bar{\omega}$ is the maximizer of $|\Psi(\omega)|$.

The same argument applies if we choose two *good* θ -shades independently and uniformly. More precisely, we claim that

$$\mathbb{E}[\omega(\xi_1, \xi_2) | \xi_1, \xi_2 \text{ are good}] = \bar{\omega}. \quad (16)$$

The reason for this is that any good θ -shade is judicious and therefore has the following property.

For each clause type $\ell = (t^1, \dots, t^h)$, any $j \in [h]$ and any color $c \in \{\mathbf{r}, \mathbf{b}, \mathbf{g}, \mathbf{y}\}$, a t_c^j fraction of clauses of type ℓ contain a clone colored c in their j th position. (17)

Given that ξ_1, ξ_2 are good, they both satisfy (17). But if we fix a clause type ℓ , an index j and a color c , then the actual *sets* of clauses of type ℓ that contain color c in their j th position are chosen independently for ξ_1, ξ_2 . This implies (16). Furthermore, again by standard arguments $\omega(\xi_1, \xi_2)$ is tightly concentrated about its conditional expectation. Hence, $\bar{\omega}$ is the maximizer of $|G(\omega)|$. $\square$

Finally, Lemma 3 implies easily that $\bar{\mathcal{O}}$ is a stationary point of $\mathcal{P}(\mathcal{O})$. The reason for this is that if ξ_1, ξ_2 are θ -shades with overlap $\bar{\mathcal{O}}$, then $\mathbb{E}[\omega(\xi_1, \xi_2)] = \bar{\omega}$, and $\omega(\xi_1, \xi_2)$ is tightly concentrated about this value (say, by Azuma’s inequality). Thus, we have established **SMM1**.

While thus far we managed to avoid an explicit expansion of the functions $\mathcal{N}(\mathcal{O})$, $\mathcal{P}(\omega)$, this becomes inevitable to prove **SMM2**. The resulting formulas are intricate due to the large number of parameters involved, but it is possible to estimate the second differentials sufficiently well to establish **SMM2**. Furthermore, step **SMM3** is largely based on combinatorial arguments again. These hinge on the notion of separability. Roughly speaking, separability ensures that we only need to consider $\mathcal{O}$ such that $\|\mathcal{O} - \bar{\mathcal{O}}\|_2 \leq O_k(k^4 2^{-k/2})$, which substantially confines the space over which we need to optimize.

Acknowledgment. This work has benefited from conversations with Dimitris Achlioptas, Florent Krzakala, Guilhem Semerjian and Lenka Zdeborová. I am also grateful to Charilaos Efthymiou for his comments on a draft version of this paper. Last but by no means least, I thank Konstantinos Panagiotou for many deep, helpful and inspiring discussions.

6. REFERENCES

- [1] D. Achlioptas, A. Coja-Oghlan: Algorithmic barriers from phase transitions. Proc. 49th FOCS (2008) 793–802.
- [2] D. Achlioptas, C. Moore: Random k -SAT: two moments suffice to cross a sharp threshold. SIAM Journal on Computing **36** (2006) 740–762.
- [3] D. Achlioptas, Y. Peres: The threshold for random k -SAT is $2^k \ln 2 - O(k)$. Journal of the AMS **17** (2004) 947–973.
- [4] D. Achlioptas, F. Ricci-Tersenghi: Random formulas have frozen variables. SIAM J. Comput. **39** (2009) 260–280.
- [5] A. Braunstein, M. Mézard, R. Zecchina: Survey propagation: an algorithm for satisfiability. Random Structures and Algorithms **27** (2005) 201–226.
- [6] P. Cheeseman, B. Kanefsky, W. Taylor: Where the *really* hard problems are. Proc. IJCAI (1991) 331–337.
- [7] V. Chvátal, B. Reed: Mick gets some (the odds are on his side). Proc. 33th FOCS (1992) 620–627.
- [8] A. Coja-Oghlan: A better algorithm for random k -SAT. SIAM J. Computing **39** (2010) 2823–2864.

- [9] A. Coja-Oghlan: On belief propagation guided decimation for random k -SAT. Proc. 22nd SODA (2011) 957–966.
- [10] A. Coja-Oghlan, K. Panagiotou: Catching the k -NAESAT threshold. Proc. 44th STOC (2012) 899–908.
- [11] A. Coja-Oghlan, K. Panagiotou: Going after the k -SAT threshold. Proc. 45th STOC (2013) 705–714.
- [12] A. Coja-Oghlan, D. Vilenchik: Chasing the k -colorability threshold. Proc. 54th FOCS (2013) 380–389.
- [13] A. Coja-Oghlan, L. Zdeborová: The condensation transition in random hypergraph 2-coloring. Proc. 23rd SODA (2012) 241–250.
- [14] J. Ding, A. Sly, N. Sun: Satisfiability threshold for random regular NAE-SAT. arXiv:1310.4784 (2013).
- [15] J. Ding, A. Sly, N. Sun: Maximum independent sets on random regular graphs. arXiv:1310.4787 (2013).
- [16] O. Dubois, Y. Boufkhad: A general upper bound for the satisfiability threshold of random r -SAT formulae. J. Algorithms **24** (1997) 395–420.
- [17] O. Dubois, J. Mandler: The 3-XORSAT threshold. Proc. 43rd FOCS (2002) 769–778.
- [18] E. Friedgut: Sharp thresholds of graph properties, and the k -SAT problem. J. AMS **12** (1999) 1017–1054.
- [19] A. Frieze, S. Suen: Analysis of two simple heuristics on a random instance of k -SAT. Journal of Algorithms **20** (1996) 312–355.
- [20] A. Frieze, N. Wormald: Random k -Sat: a tight threshold for moderately growing k . Combinatorica **25** (2005) 297–305.
- [21] A. Goerdt: A threshold for unsatisfiability. Proc. 17th MFCS (1992) 264–274.
- [22] A. Kaporis, L. Kirousis, E. Lalas: The probabilistic analysis of a greedy satisfiability algorithm. Random Structures and Algorithms **28** (2006) 444–480.
- [23] S. Kirkpatrick, B. Selman: Critical behavior in the satisfiability of random boolean expressions. Science **264** (1994) 1297–1301.
- [24] L. Kirousis, E. Kranakis, D. Krizanc, Y. Stamatiou: Approximating the unsatisfiability threshold of random formulas. Random Structures Algorithms **12** (1998) 253–269.
- [25] L. Kroc, A. Sabharwal, B. Selman: Message-passing and local heuristics as decimation strategies for satisfiability. Proc 24th SAC (2009) 1408–1414.
- [26] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, L. Zdeborová: Gibbs states and the set of solutions of random constraint satisfaction problems. Proc. National Academy of Sciences **104** (2007) 10318–10323.
- [27] E. Maneva, E. Mossel, M. Wainwright: A new look at survey propagation and its generalizations. J. ACM **54** (2007).
- [28] E. Maneva, A. Sinclair: On the satisfiability threshold and clustering of solutions of random 3-SAT formulas. Theoretical Computer Science **407** (2008) 359–369.
- [29] C. McDiarmid, Concentration for independent permutations, Combinatorics, Probability and Computing (2002) **11**, 163–178.
- [30] S. Mertens, M. Mézard, R. Zecchina: Threshold values of random K -SAT from the cavity method. Random Struct. Alg. **28** (2006) 340–373.
- [31] M. Mézard, A. Montanari: Information, physics and computation. Oxford University Press 2009.
- [32] M. Mézard, G. Parisi, R. Zecchina: Analytic and algorithmic solution of random satisfiability problems. Science **297** (2002) 812–815.
- [33] M. Molloy: The freezing threshold for k -colourings of a random graph. Proc. 43rd STOC (2012) 921–930.
- [34] R. Monasson, R. Zecchina: Entropy of the K -satisfiability problem. Physical Review Letters **76** (1996) 3881–3885.
- [35] B. Pittel, G. Sorkin: The satisfiability threshold for k -XORSAT. arXiv:1212.1905 (2012).
- [36] V. Rathi, E. Aurell, L. K. Rasmussen, M. Skoglund: Bounds on threshold of regular random k -SAT. Proc. 12th SAT (2010) 264–277.