

# Catching the $k$ -NAESAT Threshold

[Extended Abstract]

Amin Coja-Oghlan<sup>\*</sup>  
University of Warwick  
Zeeman building  
Coventry CV4 7AL, UK  
a.coja-oghlan@warwick.ac.uk

Konstantinos Panagiotou<sup>†</sup>  
Max Planck Institute for Informatics  
Campus E1.4  
66123 Saarbrücken, Germany  
kpanagio@mpi-inf.mpg.de

## ABSTRACT

The best current estimates of the thresholds for the existence of solutions in random constraint satisfaction problems (‘CSPs’) mostly derive from the *first* and the *second moment method*. Yet apart from a very few exceptional cases these methods do not quite yield matching upper and lower bounds. According to deep but non-rigorous arguments from statistical mechanics, this discrepancy is due to a change in the geometry of the set of solutions called *condensation* that occurs shortly before the actual threshold for the existence of solutions (Krzakala, Montanari, Ricci-Tersenghi, Semerjian, Zdeborová: PNAS 2007). To cope with condensation, physicists have developed a sophisticated but non-rigorous formalism called *Survey Propagation* (Mézard, Parisi, Zecchina: Science 2002). This formalism yields precise conjectures on the threshold values of many random CSPs. Here we develop a new *Survey Propagation inspired second moment method* for the random  $k$ -NAESAT problem, which is one of the standard benchmark problems in the theory of random CSPs. This new technique allows us to overcome the barrier posed by condensation rigorously. We prove that the threshold for the existence of solutions in random  $k$ -NAESAT is  $2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) + \varepsilon_k$ , where  $|\varepsilon_k| \leq 2^{-(1-o_k(1))k}$ , thereby verifying the statistical mechanics conjecture for this problem.

## Categories and Subject Descriptors

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

## General Terms

Theory

<sup>\*</sup>Supported by EPSRC grant EP/G039070/2 and ERC Starting Grant 278857-PTCC (FP7).

<sup>†</sup>Parts of this work were performed while the second author was visiting University of Warwick.

Permission to make digital or hard copies of all or part 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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC’12, May 19–22, 2012, New York, New York, USA.  
Copyright 2012 ACM 978-1-4503-1245-5/12/05 ...\$10.00.

## Keywords

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

## 1. INTRODUCTION

Over the past decade, physicists have developed sophisticated but non-rigorous techniques for the study of random constraint satisfaction problems (‘CSPs’) such as random  $k$ -SAT or random graph  $k$ -coloring [27, 29]. This work has led to a remarkably detailed *conjectured* picture, according to which various phase transitions affect both the combinatorial and computational nature of random problems. By now, some of these predictions have been turned into rigorous theorems. Examples include results on the ‘shattering’ of the solution space [1, 7], work on (non-)reconstruction and sampling [18, 24, 30], and even new algorithms for random CSPs [9, 19]. Many of these contributions have led to the development of new rigorous techniques. Indeed, it seems fair to say that, combined, these results have advanced our understanding of random CSPs quite significantly.

However, thus far substantial bits of the statistical mechanics picture have eluded all rigorous attempts. Perhaps most importantly, apart from a very few special cases, the thresholds for the existence of solutions in random CSPs have not been pinned down exactly. While rigorous upper and lower bounds can be derived via the *first* and the *second moment method* [5], these bounds do not quite match in most examples, including prominent ones such as random  $k$ -SAT or random graph  $k$ -coloring. In fact, the statistical mechanics techniques suggest a striking explanation for this discrepancy, namely the existence of a *condensation phase* shortly before the threshold for the existence of solutions. In this phase, a crucial necessary condition for the success of the (standard) second moment method is violated. Indeed, in statistical mechanics a deep formalism called *Survey Propagation* (‘SP’) has been developed expressly to deal with condensation. While SP is primarily an analysis technique, a spin-off has been the *SP guided decimation* algorithm, which seems highly successful at solving random CSPs experimentally.

In this paper we propose a new *SP-inspired second moment method* that allows us to overcome the barrier posed by condensation. The specific problem that we work with is random  $k$ -NAESAT, one of the standard benchmark problems in the theory of random CSPs. Random  $k$ -NAESAT is technically a bit simpler than random  $k$ -SAT due to a certain symmetry property, but computationally and struc-

turally both problems have strong similarities. We determine the threshold for the existence of solutions in random  $k$ -NAESAT up to an additive error that tends to zero exponentially with  $k$ . This is the first time that the threshold in any random CSP of this type can be calculated with such accuracy. While from a technical viewpoint  $k$ -NAESAT is perhaps the simplest example of a random CSP that exhibits condensation, our proof technique rests on a rather generic approach. Therefore, we believe that with additional technical work our approach can be extended to many other problems, including random  $k$ -SAT or random graph  $k$ -coloring.

To define random  $k$ -NAESAT formally, let  $k \geq 3$  and  $n > 0$  be integers and let  $V = \{x_1, \dots, x_n\}$  be a set of Boolean variables. For a fixed real  $r > 0$  we let  $m = m(n) = \lceil rn \rceil$ . Further, let  $\Phi = \Phi_k(n, m)$  be a propositional formula obtained by choosing  $m$  clauses of length  $k$  over  $V$  uniformly and independently at random among all  $(2n)^k$  possible clauses. We say that an assignment  $\sigma : V \rightarrow \{0, 1\}$  is an *NAE-solution* (a ‘‘solution’’) if each clause has both a literal that evaluates to ‘true’ under  $\sigma$  and one that evaluates to ‘false’. In other words, both  $\sigma$  and its inverse  $\bar{\sigma} : x_i \mapsto 1 - \sigma(x_i)$  are satisfying assignments of the Boolean formula  $\Phi$ . We say that an event occurs *with high probability* (‘‘w.h.p.’’) if its probability tends to one as  $n \rightarrow \infty$ .

Friedgut [22] proved that for any  $k$  there exists a *sharp threshold sequence*  $r_{k\text{-NAE}} = r_{k\text{-NAE}}(n)$  such that for any fixed  $\varepsilon > 0$  w.h.p.  $\Phi$  has a NAE-solution if  $r < r_{k\text{-NAE}} - \varepsilon$ , while w.h.p.  $\Phi$  fails to have one if  $r > r_{k\text{-NAE}} + \varepsilon$ . It is widely conjectured but as yet unproven that the threshold sequence converges for any  $k \geq 3$ . The best previous bounds on  $r_{k\text{-NAE}}$  were derived by Achlioptas and Moore [3] and Coja-Oghlan and Zdeborova [12] via the first/second moment method:

$$\begin{aligned} r_{\text{second}} &= 2^{k-1} \ln 2 - \ln 2 + o_k(1) \\ &\leq r_{k\text{-NAE}} \\ &\leq r_{\text{first}} = 2^{k-1} \ln 2 - \frac{\ln 2}{2} + o_k(1), \end{aligned} \quad (1)$$

where  $o_k(1)$  hides a term that tends to 0 for large  $k$ . This left an additive gap of  $\frac{1}{2} \ln 2 \approx 0.347$ , which our main result closes.

**THEOREM 1.** *There is a sequence  $\varepsilon_k = 2^{-(1-o_k(1))k}$  such that*

$$\left| r_{k\text{-NAE}} - \left( 2^{k-1} \ln 2 - \left( \frac{\ln 2}{2} + \frac{1}{4} \right) \right) \right| \leq \varepsilon_k. \quad (2)$$

While the numerical improvement obtained in Theorem 1 may seem modest, we are going to argue that the result is conceptually quite significant for two reasons. First, we obtain (virtually) matching upper and lower bounds for the first time in a random CSP of this type. Second, and perhaps even more importantly, we devise a rigorous method for taming the condensation phenomenon. Indeed, condensation has been the main obstacle to determining the precise thresholds in random CSPs for the past decade. To understand why, we need to discuss the statistical mechanics picture and its relation to the second moment method.

## 2. CONDENSATION AND THE SECOND MOMENT METHOD

*The statistical mechanics perspective.* We follow [27] to sketch the non-rigorous statistical mechanics approach on random  $k$ -NAESAT. Let  $\mathcal{S}(\Phi) \subset \{0, 1\}^n$  denote the set of NAE-solutions of  $\Phi$ , and let  $Z(\Phi) = |\mathcal{S}(\Phi)|$  be the number of solutions. We turn  $\mathcal{S}(\Phi)$  into a graph by considering two solutions  $\sigma, \tau$  adjacent if their Hamming distance is  $o(n)$ . According to [27], the ‘shape’ of  $\mathcal{S}(\Phi)$  undergoes two substantial changes w.h.p. at certain densities  $0 < r_{\text{sh}} < r_{\text{cond}} < r_{k\text{-NAE}}$ .

The first transition occurs at  $r_{\text{sh}} \sim 2^{k-1} \ln(k)/k$ , almost a factor of  $k$  below  $r_{k\text{-NAE}}$ . Namely, for  $r < r_{\text{sh}}$ ,  $\mathcal{S}(\Phi)$  is (essentially) a connected graph. But in the *shattering phase*  $r_{\text{sh}} < r < r_{\text{cond}}$ ,  $\mathcal{S}(\Phi)$  splits into connected components  $S_1, \dots, S_{N(\Phi)}$  called *clusters* that are mutually separated by a linear Hamming distance  $\Omega(n)$ . Each cluster  $S_i$  only comprises an exponentially small fraction of  $\mathcal{S}(\Phi)$ . In particular, the total number  $N(\Phi)$  of clusters, the so-called *complexity*, is exponential in  $n$ . This ‘shattering’ of  $\mathcal{S}(\Phi)$  was indeed established rigorously in [1].

As the density  $r$  increases beyond  $r_{\text{sh}}$ , both the overall number  $Z(\Phi)$  of solutions and the number and sizes of the clusters shrink. However, the cluster sizes decrease at a slower rate than  $Z(\Phi)$ , until at density  $r_{\text{cond}} = 2^{k-1} \ln 2 - \ln 2 + o_k(1)$  the largest cluster has size essentially  $\Omega(Z(\Phi))$  w.h.p., and thus constitutes a significant fraction of the entire solution space. In effect, in the *condensation phase*  $r_{\text{cond}} < r < r_{k\text{-NAE}}$ , the set  $\mathcal{S}(\Phi)$  still decomposes into an exponential number of clusters  $S_1, \dots, S_{N(\Phi)}$ , each of tiny diameter and all mutually separated by Hamming distance  $\Omega(n)$ . But in contrast to the shattered phase, now the largest cluster contains a *constant* fraction of the entire set  $\mathcal{S}(\Phi)$ . Indeed, w.h.p. a *bounded* number of clusters contain a  $1 - o(1)$ -fraction of all solutions.

The dominance of a few large clusters in the condensation phase complicates the probabilistic nature of the problem dramatically. To see why, consider the experiment of first choosing a random formula  $\Phi$ , and then picking two solutions  $\sigma, \tau \in \mathcal{S}(\Phi)$  uniformly and independently. For  $r_{\text{sh}} < r < r_{\text{cond}}$ ,  $\sigma, \tau$  likely belong to different clusters, and hence can be expected to have a ‘large’ Hamming distance. In fact, it is implicit in the previous work on the second moment method that  $\text{dist}(\sigma, \tau) \sim n/2$  w.h.p. [3, 12]. Intuitively, this means that the two random solutions ‘decorrelate’. By contrast, for  $r_{\text{cond}} < r < r_{k\text{-NAE}}$  both  $\sigma, \tau$  belong to the same large cluster with a non-vanishing probability. In effect, with a non-vanishing probability their distance  $\text{dist}(\sigma, \tau)$  is tiny, reflecting that solutions in the same cluster are heavily correlated.

The purpose of the physicists’ *Survey Propagation* technique is precisely to deal with this type of correlation. The basic idea is to work with a different, non-uniform probability distribution on  $\mathcal{S}(\Phi)$ . This *SP distribution* is induced by first choosing a *cluster*  $S_i$  uniformly at random among  $S_1, \dots, S_{N(\Phi)}$ , and then selecting a solution in that cluster  $S_i$  uniformly. Since the number  $N(\Phi)$  of clusters is (thought to be) exponential in  $n$  throughout the condensation phase, two solutions  $\sigma', \tau'$  chosen independently from the SP distribution are expected to lie in distinct clusters and thus to decorrelate w.h.p.

Starting from this (appropriately formalized) decorrela-

tion assumption, the SP formalism prescribes a sequence of delicate (non-rigorous) steps to reduce the computation of the *precise* threshold  $r_{k\text{-NAE}}$  to the solution of a continuous variational problem for any  $k \geq 3$  [14, 31]. This variational problem is itself highly non-trivial, but heuristic numerical techniques yield plausible approximations for small values of  $k$  [28]. Moreover, asymptotically for large  $k$  the variational problem can be solved analytically. This led to the conjecture that  $r_{k\text{-NAE}} = 2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) + o_k(1)$  [14], which Theorem 1 resolves.

Is Theorem 1 “optimal”? Of course, it would be interesting to prove that for any  $k$ , the *precise* threshold  $r_{k\text{-NAE}}$  equals the solution to the variational problem that the SP formalism spits out. However, given that this continuous problem itself appears difficult to solve analytically (to say the very least), it seems that such a result would merely establish the equivalence of two hard mathematical problems. Thus, we believe that Theorem 1 marks the end of the line as far as an analytic/explicit computation of  $r_{k\text{-NAE}}$  is concerned. However, our results strongly suggest that these two problems really are equivalent, at least in the limit of large  $k$ .

**The first and the second moment method.** The above statistical mechanics picture holds the key to understanding why the previous arguments did not suffice to pin down  $r_{k\text{-NAE}}$  precisely. The best previous bounds (1) were obtained by applying the first/second moment method to the number  $Z(\Phi)$  of solutions, or a closely related random variable.

With respect to the upper bound, if for some density  $r$  the first moment  $E[Z(\Phi)]$  tends to 0 as  $n$  gets large, then  $Z(\Phi) = 0$  w.h.p. by Markov’s inequality. Thus,  $r_{k\text{-NAE}} \leq r$ . Indeed, it is not difficult to verify that  $E[Z(\Phi)] = o(1)$  for  $r = r_{\text{first}}$  [3]. This gives the upper bound in (1).

The purpose of the second moment method is to bound  $r_{k\text{-NAE}}$  from below. The general approach is this: suppose we can define a random variable  $Y = Y(\Phi) \geq 0$  such that  $Y > 0$  only if  $\Phi$  has a NAE-solution. Moreover, assume that for some density  $r$ , the second moment  $E[Y^2]$  satisfies

$$E[Y^2] \leq C \cdot E[Y]^2 \quad (3)$$

with  $C = C(k) \geq 1$  dependent on  $k$  but not on  $n$ . Then the *Paley-Zygmund inequality*  $P[Y > 0] \geq E[Y]^2 / E[Y^2]$  implies that

$$\begin{aligned} P[\Phi \text{ has a NAE-solution}] &\geq P[Y > 0] \\ &\geq E[Y^2] / E[Y]^2 \geq 1/C > 0. \end{aligned} \quad (4)$$

Because the  $k$ -NAESAT threshold is sharp, and as  $C$  is independent of  $n$ , (4) implies that  $r_{k\text{-NAE}} \geq r$ .

The obvious choice of random variable is the number  $Z(\Phi)$  of solutions. Since  $Z(\Phi)^2$  is just the number of *pairs* of NAE-solutions, the second moment can be written as

$$E[Z(\Phi)^2] = \sum_{\sigma, \tau \in \{0,1\}^n} P[\text{both } \sigma, \tau \text{ are NAE-solutions}]. \quad (5)$$

Indeed, Achlioptas and Moore [3] proved that (3) is satisfied for  $Y = Z(\Phi)$  if  $r \leq 2^{k-1} \ln 2 - (1 + \ln 2)/2$ . Improving upon [3], Coja-Oghlan and Zdeborová [12] obtained the best previous lower bound (1) by considering a slightly modified random variable  $Z'(\Phi)$ . Namely,  $Z'(\Phi) = Z(\Phi) \cdot \mathbf{1}_{\Phi \in \mathcal{A}}$ , where  $\mathcal{A}$  is a certain event such that  $\Phi \in \mathcal{A}$  w.h.p. In other words,  $Z'(\Phi)$  is equal to  $Z(\Phi)$  for almost all formulas, but a small fraction of “bad” formulas (that would blow up the

second moment) are excluded. Still,  $Z'(\Phi)$  admits a similar decomposition as (5) (one just has to condition on  $\mathcal{A}$ ).

As (5) shows, the second moment analysis of either  $Z(\Phi)$  or  $Z'(\Phi)$  boils down to studying the correlations amongst *pairs* of solutions. In fact, it was observed in [3, 12] that a *necessary* condition for the success of this approach is that two independently and uniformly chosen  $\sigma, \tau \in \mathcal{S}(\Phi)$  satisfy  $\text{dist}(\sigma, \tau) \sim n/2$  w.h.p. But according to the statistical mechanics picture, this decorrelation condition is violated for  $r > r_{\text{cond}}$  due to the presence of large clusters. Therefore, it is not surprising that the best previous lower bound (1) on  $r_{k\text{-NAE}}$  coincides with the (conjectured) condensation threshold  $r_{\text{cond}}$ . Indeed, it was verified in [12] that a certain “weak” form of condensation sets in at  $r \sim r_{\text{cond}}$ .

The statistical mechanics prescription to overcome these correlations is to work with the Survey Propagation distribution (first select a cluster uniformly, then choose a random solution from that cluster) rather than the uniform distribution over  $\mathcal{S}(\Phi)$ . This is precisely the key idea behind our new *SP-inspired second moment argument*. Roughly speaking, we are going to develop a way to apply the second moment method to the number  $N(\Phi)$  of *clusters*, rather than the number of solutions. More precisely, we introduce a parameter  $\beta$  that allows us to work with clusters of a prescribed size. A specific choice of  $\beta$  (namely,  $\beta = 1/2$ ) corresponds to the SP distribution and thus to working with  $Y(\Phi) = N(\Phi)$ .

This new technique allows us to obtain various further results. For instance, we can pin down the typical values of both  $Z(\Phi)$  and  $N(\Phi)$  throughout the condensation phase (details omitted). Furthermore, our proof entails the following result that confirms the physics conjecture that pairs of solutions drawn from the SP distribution decorrelate throughout the condensation phase.

COROLLARY 1. *Suppose that*

$$r_{\text{cond}} \leq r \leq 2^{k-1} \ln 2 - \left( \frac{\ln 2}{2} + \frac{1}{4} \right) - \varepsilon_k.$$

*Let  $\sigma', \tau'$  be drawn independently from the SP distribution. Then  $\text{dist}(\sigma', \tau') = (\frac{1}{2} + o_k(1))n$  w.h.p.*

### 3. RELATED WORK

**Rigorous work.** The  $k$ -NAESAT problem is well-known to be NP-complete in the worst case for any  $k \geq 3$ . In fact, the NP-complete problem of 2-coloring a  $k$ -uniform hypergraph (with  $k \geq 3$ ) simply is the special case of  $k$ -NAESAT without negations. The results in [12] are actually phrased in terms of hypergraph 2-coloring but carry over to  $k$ -NAESAT directly.

The main contribution of Theorem 1 is the improved *lower* bound. In fact, the upper bound in (2) can be obtained in several different ways. Achlioptas and Moore [3] state without proof that the (quite intricate) enhanced first moment argument from [16, 26] can be used to show that  $r_{k\text{-NAE}} \leq 2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) + o_k(1)$ . This is indeed plausible as, in terms of the statistical mechanics intuition (which was unknown to the authors of [16, 26]) this argument amounts to computing the first moment of the number of *clusters*. Alternatively, generalizing work of Franz and Leone [21], Panchenko and Talagrand [31] proved that the variational problem that results from the SP formalism yields a rigorous upper bound on  $r_{k\text{-NAE}}$ , which is conjectured to be tight for any  $k \geq 3$ . The variational problem can

be solved asymptotically in the large- $k$  limit (unpublished), yielding the upper bound stated in Theorem 1. In this paper we obtain the upper bound by a relatively simple third argument that has a neat combinatorial interpretation.

The proofs of the lower bounds in [3, 12] and in the present paper are non-constructive in the sense that they do not entail an efficient algorithm for finding a NAE-solution w.h.p. The best current algorithm for random  $k$ -NAESAT is known to succeed for  $r \leq O_k(2^k/k)$ , thus being a factor of  $\Omega_k(k)$  below  $r_{k\text{-NAE}}$  [2].

>From a statistical mechanics point of view, many random CSPs are similar to random  $k$ -NAESAT. In particular, the physics methods suggest the existence of a condensation phase in most random CSPs (e.g., random  $k$ -SAT/graph  $k$ -coloring). While [3] provided the prototype for the second moment arguments in these and other problems, the technical details in random graph  $k$ -coloring [4] or random  $k$ -SAT [6] are quite a bit more intricate than in random  $k$ -NAESAT.

For instance, random  $k$ -NAESAT is simpler than random  $k$ -SAT because for any NAE-solution  $\sigma$  the inverse  $\bar{\sigma} : x \mapsto 1 - \sigma(x)$  is a NAE-solution as well. This symmetry of the solution space under inversion simplifies the second moment calculations significantly. To cope with the absence of symmetry in random  $k$ -SAT, Achlioptas and Peres [6] weighted satisfying assignments cleverly in order to recover the beneficial analytic properties that symmetry induces. Our new second moment method is quite different from this weighting approach, since the asymmetry that called for the weighting scheme in [6] is absent in  $k$ -NAESAT.

None of the (few) random CSPs in which the threshold for the existence of solutions is known precisely has a condensation phase. The most prominent example is random  $k$ -XORSAT (random linear equations mod 2) [17, 32]. In this case, the algebraic nature of the problem precludes condensation: all clusters are simply translations of the kernel. Similarly, the condensation phase is empty in the uniquely extendible problem from [13]. Also in random  $k$ -SAT with  $k = k(n) > \log_2 n$  (i.e., the clause length grows as a function of  $n$ ), where the precise threshold has been determined by Frieze and Wormald [23] via the second moment method, condensation does not occur [11]. Nor does it in random 2-SAT [8, 25].

Parts of our proof require a precise analysis of geometry of the solution space  $\mathcal{S}(\Phi)$ . This analysis harnesses some of the ideas that were developed in previous work [1, 7, 12, 15] (e.g., arguments for proving the existence of clusters or of “rigid variables”). However, we need to go beyond these previous arguments significantly in two respects. First, we need to generalize them to accommodate the parameter  $\beta$  that controls the cluster sizes. Second, we need rather precise quantitative information about the cluster structures.

**Survey Propagation guided decimation.** The SP formalism has given rise to an efficient message passing algorithm called *Survey Propagation guided decimation* (‘SPD’) [29]. Experimentally, SPD seems spectacularly successful at solving, e.g., random  $k$ -SAT for small values of  $k$ . Unfortunately, no quantitative analysis of this algorithm is currently known (not even a non-rigorous one). The basic idea behind SPD is to approximate the marginals of the SP distribution (i.e., the probability that a given variable is ‘true’ in a solution drawn from the SP distribution) via a message passing heuristic. Then a variable  $x$  is selected according to some rule and

is assigned a value based on the (approximate) marginal. The entire procedure is repeated on the “decimated” problem instance where  $x$  has been eliminated, until (hopefully) a solution is found.

The decorrelation of random solutions chosen from the SP distribution is a crucial assumption behind the message passing computation of the SP marginals. Corollary 1 establishes such a decorrelation property rigorously. However, in order to actually analyze SPD, one would have to generalize Corollary 1 to the situation of a “decimated” random formula in which a number of variables have already been eliminated by previous steps of the algorithm. Still, we believe that the techniques developed in this paper are a (necessary) first step towards a rigorous analysis of SPD.

## 4. HEAVY SOLUTIONS AND THE FIRST MOMENT

*In the rest of the paper we sketch the SP-inspired second moment method on which the proof of Theorem 1 is based. Aiming for an asymptotic result, we may assume that  $k \geq k_0$  for some (large) constant  $k_0 > 3$ . We also assume  $r = 2^{k-1} \ln 2 - \rho$  for some  $\frac{1}{2} \ln 2 \leq \rho \leq \ln 2$ . Let  $\Phi_i$  denote the  $i$ th clause of the random formula  $\Phi$  so that  $\Phi = \Phi_1 \wedge \dots \wedge \Phi_m$ . Furthermore, let  $\Phi_{ij}$  signify the  $j$ th literal of clause  $\Phi_i$ ; thus,  $\Phi_i = \Phi_{i1} \vee \dots \vee \Phi_{ik}$ . For a literal  $\ell$  we let  $|\ell|$  denote the underlying variable.*

As we discussed earlier, the demise of the “standard” second moment method in the condensation phase is due to the dominance of few large clusters. The statistical mechanics prescription for circumventing this issue is to work with a non-uniform distribution over solutions that favors “small” clusters. To implement this strategy, we are going to exhibit a simple parameter that governs the size of the cluster that a solution belongs to. Formally, we define the *cluster* of  $\sigma \in \mathcal{S}(\Phi)$  as

$$\mathcal{C}(\sigma) = \mathcal{C}_\Phi(\sigma) = \{\tau \in \mathcal{S}(\Phi) : \text{dist}(\sigma, \tau) \leq 0.01n\}.$$

This definition is vindicated by the following observation from [12], which shows that any two solutions either have the same cluster or are well-separated.

**PROPOSITION 1.** *Assume  $2^{k-1} \ln 2 - \ln 2 \leq r \leq r_{k\text{-NAE}}$ . W.h.p. any two  $\sigma, \tau \in \mathcal{S}(\Phi)$  either satisfy*

$$\text{dist}(\sigma, \tau) \leq 0.01n$$

or

$$\text{dist}(\sigma, \tau) \geq (1/2 - 2^{-k/3})n.$$

To proceed, we need to get an idea of the “shape” of the clusters  $\mathcal{C}(\sigma)$ . According to the SP formalism, each cluster has a set  $\mathcal{R}(\sigma)$  of  $\Omega(n)$  *rigid variables* on which *all* assignments in  $\mathcal{C}(\sigma)$  coincide, while the values of the non-rigid variables vary. Formally, we have  $\tau(x) = \sigma(x)$  for all  $x \in \mathcal{R}(\sigma)$  and all  $\tau \in \mathcal{C}(\sigma)$ , while for each  $x \notin \mathcal{R}(\sigma)$  there is  $\tau \in \mathcal{C}(\sigma)$  such that  $\tau(x) \neq \sigma(x)$ . This implies an immediate bound on the size of  $\mathcal{C}(\sigma)$ , namely  $|\mathcal{C}(\sigma)| \leq 2^{n-|\mathcal{R}(\sigma)|}$ . Indeed, we are going to prove that every cluster has a rigid set of size  $\Omega(n)$  w.h.p., and that for all clusters w.h.p.

$$\log_2 |\mathcal{C}(\sigma)| = (1 - o_k(1))(n - |\mathcal{R}(\sigma)|). \quad (6)$$

With  $|\mathcal{C}(\sigma)|$  controlled by the number of rigid variables, it might seem promising to perform first/second moment

arguments for the number of solutions with a suitably chosen number of rigid variables. The problem with this is that there is no simple way to tell whether a given variable is rigid: deciding this is NP-hard in the worst case. Intuitively, this is because rigidity emerges from the “global” interplay of variables and clauses. In effect, parametrizing by the number of rigid variables appears technically infeasible.

Instead, we are going to work with a simple “local” parameter that turns out to be a good substitute. Suppose that  $x \in \mathcal{R}(\sigma)$ . Then  $x$  must occur in some clause  $\Phi_i$  that would be violated if  $x$  was assigned the opposite value  $1 - \sigma(x)$  (with all other variables unchanged). By the definition of  $k$ -NAESAT, this means that the other  $k - 1$  literals of  $\Phi_i$  take the opposite value of the literal whose underlying variable  $x$  is. In this case we say that  $x$  *supports*  $\Phi_i$  under  $\sigma$ , and we call  $\Phi_i$  a *critical* clause. Moreover, we call a variable that supports a clause *blocked*, while all other variables are *free*. While every rigid variable is blocked, the converse is not generally true. Nonetheless, we will see that the number of variables that are blocked but not rigid is small enough so that we can control the cluster sizes in terms of blocked variables.

As a first step, we are going to estimate the expected number of solutions with a given number of blocked variables. Let  $\lambda = \frac{kr}{2^{k-1}-1} = k \ln 2 + O_k(k/2^k)$  and let us say that  $\sigma \in \mathcal{S}(\Phi)$  is  $\beta$ -heavy if exactly  $(1 - \beta) \exp(-\lambda)n$  variables are free. Let  $\mathcal{S}_\beta(\Phi)$  be the set of all  $\beta$ -heavy solutions and let  $Z_\beta = |\mathcal{S}_\beta(\Phi)|$  denote their number.

PROPOSITION 2. *For any  $\beta \leq 1$  we have*

$$\begin{aligned} \ln E[Z_\beta] &= \frac{n}{2^k} \left( 2\rho - \ln(2) - (1 - \beta) \ln(1 - \beta) - \beta + O_k(k 2^{-k}) \right). \end{aligned} \quad (7)$$

In particular,  $Z_\beta = 0$  for all  $\beta < -3/2$  w.h.p.

PROOF. The computation of  $E[Z_\beta]$  is instructive because it hinges upon the solution of an occupancy problem that will play an important role in the second moment computation. Let  $\mathbf{1}$  denote the assignment that sets all variables to true. By the linearity of expectation and by symmetry, we have

$$\begin{aligned} E[Z_\beta] &= \sum_{\sigma \in \{0,1\}^n} \mathbb{P}[\sigma \text{ is a } \beta\text{-heavy solution}] \\ &= 2^n \mathbb{P}[\mathbf{1} \text{ is a } \beta\text{-heavy solution}] \\ &= 2^n \mathbb{P}[\mathbf{1} \text{ is } \beta\text{-heavy} \mid \mathbf{1} \text{ is solution}] \mathbb{P}[\mathbf{1} \text{ is solution}]. \end{aligned}$$

Clearly,  $\mathbf{1}$  is a solution iff each clause of  $\Phi$  contains both a positive and a negative literal. A random clause has this property with probability  $1 - 2^{1-k}$ . Since the  $m \sim rn$  clauses are chosen independently, we get

$$\begin{aligned} 2^n \cdot \mathbb{P}[\mathbf{1} \text{ is a solution}] &= 2^n (1 - 2^{1-k})^m \\ &= \exp \left[ \frac{n}{2^k} \left( 2\rho - \ln 2 + O_k(2^{-k}) \right) \right]. \end{aligned}$$

Working out the conditional probability that  $\mathbf{1}$  is  $\beta$ -heavy is not so straightforward. Whether  $\mathbf{1}$  is  $\beta$ -heavy depends only on the critical clauses of  $\Phi$ . Let  $X$  be their number. Given that  $\mathbf{1}$  is a solution, each clause  $\Phi_i$  is critical with probability  $k/(2^{k-1} - 1)$  independently (as there are  $2k$  ways to choose the literal signs to obtain a critical clause). Hence,

$X$  has a binomial distribution  $\text{Bin}(m, k/(2^{k-1} - 1))$  with mean

$$E[X \mid \mathbf{1} \in \mathcal{S}(H)] = \frac{km}{2^{k-1} - 1} = \lambda n.$$

Since the supporting variable of each critical clause is uniformly distributed, given  $\mathbf{1} \in \mathcal{S}(H)$  the *expected* number of clauses that each variable supports equals  $\lambda$ . Thinking of the variables as bins and of the critical clauses as balls, standard results on the occupancy problem show that the number of free variables is  $(1 + o(1)) \exp(-\lambda)n$  w.h.p. Thus,  $E[Z_\beta]$  is maximized for  $\beta = 0$ .

By contrast, values  $\beta \neq 0$  correspond to *atypical* outcomes of the occupancy problem. Values  $\beta < 0$  require an excess number of “empty bins”, while  $\beta > 0$  means that fewer bins than expected are empty. To determine the precise (exponentially small) probability of getting  $(1 - \beta) \exp(-\lambda)n$  empty bins, we need to balance large deviations of  $X$  against the probability that exactly  $(1 - \beta) \exp(-\lambda)n$  bins remain empty for a given value of  $X$ . The result of this combined large deviations analysis is the expression (7), see below. The analysis also shows that  $E[Z_\beta] = \exp(-\Omega(n))$  for  $\beta < -3/2$ , whence  $Z_\beta = 0$  w.h.p. for  $\beta < -3/2$ , as claimed.

Let, as above,  $X \sim \text{Bin}(m, k/(2^{k-1} - 1))$ . We throw  $X$  balls into  $n$  bins uniformly at random. Let  $B_i$  denote the number of bins that receive  $i$  balls. Then, we will show that for any  $-3/2 \leq \beta \leq 1$

$$\begin{aligned} n^{-1} \ln \Pr \left[ B_0 = (1 - \beta) e^{-\lambda} n \right] \\ = n^{-1} \ln \Pr \left[ \text{Bin}(n, e^{-\lambda}) = (1 - \beta) e^{-\lambda} n \right] + O_k(k 4^{-k}). \end{aligned} \quad (8)$$

The statement of the proposition follows from the fact that for any  $-3/2 \leq \beta \leq 1$

$$\begin{aligned} n^{-1} \ln \Pr \left[ \text{Bin}(n, e^{-\lambda}) = \lfloor (1 - \beta) e^{-\lambda} n \rfloor \right] \\ = - (1 - \beta) \ln(1 - \beta) - \beta + O_k(4^{-k}). \end{aligned}$$

This can be verified by simple algebraic manipulations using the exact expression for the density of the binomial distribution, which are omitted from this extended abstract.

We shall estimate the probability in the (8) by conditioning on any specific value  $x$  of  $X$ . Let  $F_i$  be the number of balls in the  $i$ th bin, and let  $P_1, \dots, P_n$  be independent Poisson distributed random variables with mean  $\lambda$ . It is well-known and easy to verify that the distribution of  $(F_1, \dots, F_n)$  is the same as the distribution of  $(P_1, \dots, P_n)$ , *conditioned on the event*  $\mathcal{A}(x) = “\sum_{1 \leq i \leq n} P_i = x”$ . So, if we denote by  $N_0$  the number of  $P_i$ ’s that are equal to 0, we infer that

$$\begin{aligned} \Pr \left[ B_0 = (1 - \beta) e^{-\lambda} n \mid X = x \right] \\ = \Pr \left[ N_0 = (1 - \beta) e^{-\lambda} n \mid \mathcal{A}(x) \right]. \end{aligned}$$

By the law of total probability this equals

$$\begin{aligned} \Pr \left[ B_0 = (1 - \beta) e^{-\lambda} n \mid X = x \right] \\ = \Pr \left[ N_0 = (1 - \beta) e^{-\lambda} n \right] \cdot \frac{\Pr[\mathcal{A}(x) \mid N_0 = (1 - \beta) e^{-\lambda} n]}{\Pr[\mathcal{A}(x)]}. \end{aligned}$$

Note that  $N_0 \sim \text{Bin}(n, e^{-\lambda})$ . Furthermore, if we denote by  $P'_1, \dots, P'_{\xi n}$ , where  $\xi = 1 - (1 - \beta) e^{-\lambda}$ , independent Poisson variables that are conditioned on being at least 1, then the

above equation implies that

$$\begin{aligned} & \frac{\Pr[B_0 = (1 - \beta)e^{-\lambda n}]}{\Pr[\text{Bin}(n, e^{-\lambda}) = (1 - \beta)e^{-\lambda n}]} \\ &= \sum_{x=\xi n}^m \frac{\Pr[\sum_{i=1}^{\xi n} P'_i = x]}{\Pr[\text{Po}(\lambda n) = x]} \cdot \Pr[\text{Bin}(rn, k/(2^{k-1} - 1)) = x]. \end{aligned} \quad (9)$$

In order to complete the proof of (8) we will derive in the sequel appropriate bounds for the right-hand side of the above equation. First, to obtain a lower bound, note that  $\xi < \lambda$ , since  $\xi < 1$  and  $\lambda = k \ln 2 + O_k(k2^{-k})$ , which is  $> 1$  for sufficiently large  $k$ . Thus, we can obtain a lower bound for (9) by considering only the term in the sum that corresponds to  $x = \lambda n$ . Since  $\mathbb{E}[\text{Po}(\lambda n)] = \mathbb{E}[\text{Bin}(rn, k/(2^{k-1} - 1))] = \lambda n$ , we infer by applying a local limit theorem that

$$\Pr[\text{Po}(\lambda n) = \lambda n] = \Theta(n^{-1/2})$$

and

$$\Pr[\text{Bin}(rn, k/(2^{k-1} - 1)) = \lambda n] = \Theta(n^{-1/2}).$$

It remains to bound  $\Pr[\sum_{i=1}^{\xi n} P'_i = \lambda n]$ . Note that  $\mathbb{E}[P'_1] = \frac{\lambda}{1 - e^{-\lambda}}$ . If we write  $N = \xi n$ , then

$$\Pr\left[\sum_{i=1}^{\xi n} P'_i = \lambda n\right] = \Pr\left[\sum_{i=1}^N P'_i = \left(\mathbb{E}[X_1] + \frac{\beta \lambda e^{-\lambda}}{\xi(1 - e^{-\lambda})}\right) N\right],$$

i.e., we require that the sum of the  $P'_i$ 's deviates from the expected value by  $O_k(k2^{-k}n)$ . Since the sums of independent random variables with finite moments have quadratic exponential tails, we conclude that the right-hand side of (9) is at least  $\exp\{-O_k(k4^{-k}n)\}$ . This shows the lower bound in (8).

In the remainder of this proof we will show an upper bound for the right-hand side of (9). To this end, we will argue that the ratio

$$\frac{\Pr[\text{Bin}(rn, k/(2^{k-1} - 1)) = \gamma \lambda n]}{\Pr[\text{Po}(\lambda n) = \gamma \lambda n]}$$

is essentially bounded for all  $x$  in the given range, from which the claim immediately follows. More specifically, let us write  $x = \gamma \lambda n$ , where  $\xi/\lambda \leq \gamma \leq r/\lambda$ . By applying Stirling's Formula  $N! = (1 + o(1))\sqrt{2\pi N}(N/e)^N$  we infer that

$$\Pr[\text{Po}(\lambda n) = \gamma \lambda n] = \Theta(1) n^{-1/2} \exp\{\lambda n(-1 + \gamma - \gamma \ln \gamma)\}. \quad (10)$$

Moreover, by abbreviating  $p = k/(2^{k-1} - 1)$  we get

$$\begin{aligned} & \Pr[\text{Bin}(rn, k/(2^{k-1} - 1)) = \gamma \lambda n] \\ &= \binom{rn}{(\gamma p) rn} p^{(\gamma p) rn} (1 - p)^{(1 - \gamma p) rn}. \end{aligned}$$

Since  $\binom{N}{\alpha N} \leq e^{H(\alpha)N}$ , where  $H$  denotes the entropy function, we obtain after some elementary algebra

$$\begin{aligned} & \Pr[\text{Bin}(rn, p) = \gamma \lambda n] \\ & \leq \exp\left\{\lambda n \left(-\gamma \ln \gamma - \frac{1 - \gamma p}{p} \ln \left(\frac{1 - \gamma p}{1 - p}\right)\right)\right\}. \end{aligned}$$

By combining this with (10) we obtain the estimate

$$\frac{\Pr[\text{Bin}(rn, k/(2^{k-1} - 1)) = \gamma \lambda n]}{\Pr[\text{Po}(\lambda n) = \gamma \lambda n]} \leq \Theta(\sqrt{n}) e^{f(\gamma) \lambda n},$$

where

$$f(\gamma) = 1 - \gamma - \frac{1 - \gamma p}{p} \ln \left(\frac{1 - \gamma p}{1 - p}\right).$$

Recall that  $0 < \xi/\lambda \leq \gamma \leq r/\lambda = 1/p$ , and note that both  $f(0)$  and  $f(1/p)$  are  $< 0$ . Moreover,  $f$  has an extremal point at  $\gamma = 1$ , where  $f(1) = 0$ . Thus, for all  $\gamma$  in the considered range we have that  $f(\gamma) \leq 0$ , which implies that the right-hand side of (9) is bounded from above by at most a polynomial in  $n$ . This completes the proof.  $\square$

As a next step, we need to estimate the cluster size of a  $\beta$ -heavy solution.

**PROPOSITION 3.** *W.h.p. for all  $-3/2 \leq \beta \leq 1$  all  $\beta$ -heavy  $\sigma \in \mathcal{S}(\Phi)$  satisfy*

$$\log_2 |\mathcal{C}(\sigma)| = \frac{n}{2^k} [1 - \beta + o_k(1)]. \quad (11)$$

**PROOF.** The crucial thing to show is that all but a very few blocked variables are rigid. The proof of this builds upon arguments developed in [1] to establish rigidity. Suppose that  $x$  is blocked in  $\sigma \in \mathcal{S}_\beta(\Phi)$ , i.e.,  $x$  supports some clause, say  $\Phi_1$ . In any solution  $\tau$  with  $\tau(x) \neq \sigma(x)$  there must be another variable  $x'$  that occurs in  $\Phi_1$  such that  $\tau(x') \neq \sigma(x')$ . Given that  $x$  supports  $\Phi_1$ , the other  $k - 1$  variables of  $\Phi_1$  are uniformly distributed. Since  $\sigma$  has no more than  $(1 - \beta) \exp(-\lambda)n = (1 - \beta + o_k(1))2^{-k}n$  free variables, the probability that  $x'$  is free is bounded by  $(1 - \beta + o_k(1))(k - 1)/2^k$ . In fact, since the *expected* number of clauses that each variable supports is  $\lambda = (1 + o_k(1))k \ln 2$ , it is quite likely that  $x'$  supports several clauses and that therefore “flipping”  $x'$  necessitates *several* further flips. Continuing this argument, we see that the number of flips follows a branching process with (initial) successor rate  $\lambda$ . A detailed analysis shows that for all but  $O_k(k4^{-k})n$  blocked initial variables  $x$  this process will lead to an avalanche of more than  $0.01n$  flips, whence  $\tau \notin \mathcal{C}(\sigma)$ . This shows that all but  $o_k(2^{-k})n$  blocked variables are rigid.  $\square$

We are ready to prove that  $r_{k-\text{NAE}} \leq 2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) + o_k(1)$ , which is (almost) the upper bound promised in Theorem 1. (Some additional technical work is needed to replace the  $o_k(1)$  by an error term that decays exponentially.) Let  $N_\beta = |\{\mathcal{C}(\sigma) : \sigma \in \mathcal{S}(\Phi) \text{ is } \beta\text{-heavy}\}|$  be the number of *clusters* centered around  $\beta$ -heavy solutions. By Proposition 3, each such cluster has size  $|\mathcal{C}(\sigma)| = 2^{n(1-\beta+o_k(1))/2^k}$  w.h.p. Hence, once more by Proposition 3, any solution  $\tau \in \mathcal{C}(\sigma)$  is  $\beta'$ -heavy for some  $\beta'$  satisfying  $|\beta' - \beta| \leq \delta_k = o_k(1)$  w.h.p. Letting  $Z_\beta^*$  be the total number of  $\beta'$ -heavy solutions with  $|\beta' - \beta| \leq \delta_k$ , we conclude that

$$N_\beta \cdot 2^{n(1-\beta+o_k(1))/2^k} \leq Z_\beta^* \quad \text{w.h.p.} \quad (12)$$

Clearly,  $Z_\beta^* \leq \mathbb{E}[Z_\beta^*] \cdot \exp(o(n))$  w.h.p. by Markov's inequality. Furthermore, as the total number of free variables in each cluster is an integer between 0 and  $n$ , we have  $\mathbb{E}[Z_\beta^*] \leq (n+1) \cdot \max_{\beta'} \mathbb{E}[Z_{\beta'}]$ . Combining these inequalities with the estimate of  $\mathbb{E}[Z_{\beta'}]$  from Proposition 2, we find

$$\begin{aligned} Z_\beta^* & \leq \exp[o(n)] \mathbb{E}[Z_\beta^*] \\ & \leq \exp\left(\frac{n}{2^k} [2\rho - \ln(2) - (1 - \beta) \ln(1 - \beta) - \beta + o_k(1)]\right) \end{aligned} \quad (13)$$

with high probability. Combining (12) and (13), we obtain

FACT 1. *W.h.p. we have  $N_\beta \leq \exp[\eta(\beta) \cdot n/2^k]$  for all  $\beta$ , with*

$$\eta(\beta) = 2\rho - \ln(2) - (1 - \beta)\ln(2 - 2\beta) - \beta + o_k(1). \quad (14)$$

Finally, it is a mere exercise in calculus to verify that at density  $r^* = 2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) + o_k(1)$  the exponent  $\eta(\beta)$  is negative for *all*  $\beta$ . Therefore, Fact 1 implies that  $r^*$  is an upper bound on  $r_{k\text{-NAE}}$ .

REMARK 1. *The exponent  $\eta(\beta)$  attains its maximum at  $\beta = \frac{1}{2} + o_k(1)$ . Together with our second moment bound below, this implies that for  $\beta = \frac{1}{2} + o_k(1)$  we have  $N(\Phi) = \exp(o_k(1)n) \cdot N_\beta(\Phi)$  w.h.p., i.e., setting  $\beta = \frac{1}{2} + o_k(1)$  corresponds to the uniform distribution over clusters and thus to the SP distribution.*

## 5. THE SECOND MOMENT

**A first attempt.** The obvious approach to proving a matching lower bound on  $r_{k\text{-NAE}}$  seems to be a second moment argument for the number  $Z_\beta$  of  $\beta$ -heavy solutions, for some suitable  $\beta$ . There is a subtle issue with this, but exploring it will put us on the right track.

We already computed  $E[Z_\beta]$  in Proposition 2. As  $E[Z_\beta^2]$  is the expected number of *pairs* of  $\beta$ -heavy solutions, the symmetry properties of the random formula  $\Phi$  imply that

$$E[Z_\beta^2] = E[Z_\beta] \cdot E[Z_\beta | \sigma \in \mathcal{S}_\beta(\Phi)] \text{ for any fixed } \sigma \in \{0, 1\}^n.$$

Thus, the second moment condition (3) that we would like to establish for  $Y = Z_\beta$  becomes

$$E[Z_\beta | \sigma \in \mathcal{S}_\beta(\Phi)] \leq C \cdot E[Z_\beta]. \quad (15)$$

What value of  $\beta$  should we go for? By Fact 1 a necessary condition for the existence of  $\beta$ -heavy solutions is that the exponent  $\eta(\beta)$  from (14) is positive. Let us call  $\beta$  *feasible* for a density  $r$  if it is. An elementary calculation shows that for  $r > r_{\text{cond}} = 2^{k-1} \ln 2 - \ln 2 + o_k(1)$ , any feasible  $\beta$  is strictly positive.

However, (15) turns out to be false for *any*  $\beta > 0$ , for any density  $r > 0$ . To understand why, let us define the *degree*  $d_x$  of a variable  $x \in V$  as the number of times that  $x$  occurs in the formula  $\Phi$ . Let  $\mathbf{d} = (d_x)_{x \in V}$  be the degree sequence of  $\Phi$ . It is well known that in the “plain” random formula  $\Phi$  (without conditioning on  $\sigma \in \mathcal{S}_\beta(\Phi)$ ), the degree of each variable is asymptotically Poisson with mean  $km/n$ . On the other hand, if we condition on  $\sigma \in \mathcal{S}_\beta(\Phi)$  for some  $\beta > 0$ , then the degrees are *not* asymptotically Poisson anymore. Indeed, the degree  $d_x$  is the sum of the number  $s_x$  of clauses that  $x$  supports, and the number  $d'_x$  of times that  $x$  appears otherwise. While  $d'_x$  is asymptotically Poisson with mean  $< km/n$  as the non-critical clauses do not affect the number of blocked variables at all,  $s_x$  is not. More precisely, we saw in the proof of Proposition 2 that for  $\beta > 0$ ,  $s_x$  is the number of “balls” that  $x$  receives in an *atypical* outcome of the occupancy problem. The precise distribution of  $s_x$  is quite non-trivial, but it is not difficult to verify that  $s_x$  does *not* have a Poisson distribution. Fleshing this observation out leads to the sobering

LEMMA 1. *For any  $\beta > 0$  and any  $r > 0$  we have*

$$E[Z_\beta | \sigma \in \mathcal{S}_\beta(\Phi)] \geq \exp(\Omega(n)) \cdot E[Z_\beta].$$

In summary, conditioning on  $\sigma \in \mathcal{S}_\beta(\Phi)$  with  $\beta > 0$  imposes a skewed degree distribution that in turn boosts the expected number of  $\beta$ -heavy solutions beyond the unconditional expectation.

**Making things work.** We tackle the issue of degree fluctuations by separating the choice of the degree sequence from the choice of the actual formula. More precisely, for a sequence  $\mathbf{d} = (d_x)_{x \in V}$  of non-negative integers such that  $\sum_{x \in V} d_x = km$  we let  $\Phi_{\mathbf{d}}$  denote a  $k$ -CNF with degree sequence  $\mathbf{d}$  chosen uniformly at random amongst all such formulas. Fixing a “typical” degree sequence  $\mathbf{d}$ , we are going to perform a second moment argument for  $\Phi_{\mathbf{d}}$ , thereby preventing fluctuations of the degrees.

How do we define “typical”? Ideally, we would like  $\mathbf{d}$  to enjoy all the properties that the degree sequence of the (unconditioned) random formula  $\Phi$  is likely to have. Formally, we let  $\mathbf{D} = \mathbf{D}_k(n, m)$  be the distribution of the degree sequence of  $\Phi$ . What we are going to show is that our second moment argument succeeds for a random degree sequence chosen from the distribution  $\mathbf{D}$  w.h.p.

DEFINITION 1. *A  $\beta$ -heavy solution  $\sigma \in \mathcal{S}(\Phi_{\mathbf{d}})$  is good if the following conditions are satisfied.*

- We have  $|\mathcal{C}(\sigma)| \leq E[Z_\beta(\Phi_{\mathbf{d}})]$ .

- There does not exist  $\tau \in \mathcal{S}(\Phi_{\mathbf{d}})$  with

$$0.01n \leq \text{dist}(\sigma, \tau) \leq (1/2 - 2^{-k/3})n.$$

- No variable supports more than  $3k$  clauses under  $\sigma$ .

The first two items mirror our analysis of the solution space from Section 4. The third one turns out to be useful for a purely technical reason.

Let  $\mathcal{S}_{g,\beta}(\Phi_{\mathbf{d}})$  be the set of good  $\beta$ -heavy solutions and set  $Z_{g,\beta}(\Phi_{\mathbf{d}}) = |\mathcal{S}_{g,\beta}(\Phi_{\mathbf{d}})|$ . We perform a second moment argument for  $Z_{g,\beta}(\Phi_{\mathbf{d}})$ , with  $\mathbf{d}$  chosen randomly from the distribution  $\mathbf{D}$ . The result is

PROPOSITION 4. *Suppose that  $\beta > 0$  is feasible. There is  $C = C(k)$  such that for a degree sequence  $\mathbf{d}$  chosen from the distribution  $\mathbf{D}$  w.h.p.  $E[Z_{g,\beta}(\Phi_{\mathbf{d}})^2] \leq C \cdot E[Z_{g,\beta}(\Phi_{\mathbf{d}})]^2$ .*

Proposition 4 shows that the second moment method for  $Z_{g,\beta}(\Phi_{\mathbf{d}})$  succeeds for feasible  $\beta$ . As we observed in Section 4, a feasible  $\beta > 0$  exists so long as  $r \leq 2^{k-1} \ln 2 - (\frac{\ln 2}{2} + \frac{1}{4}) - O_k(k^4/2^k)$ . Hence, Proposition 4 and the Paley-Zygmund inequality show that  $\Phi_{\mathbf{d}}$  is NAE-satisfiable for all such  $r$  with a non-vanishing probability for  $\mathbf{d}$  chosen randomly from  $\mathbf{D}$ . Consequently, the same is true of the unconditioned formula  $\Phi$  (because we could generate  $\Phi$  by first choosing  $\mathbf{d}$  from  $\mathbf{D}$  and then generating  $\Phi_{\mathbf{d}}$ ). Since the  $k$ -NAESAT threshold is sharp [22], we obtain the lower bound in Theorem 1.

**Proving Proposition 4.** As a first step, we need to work out  $E[Z_{g,\beta}(\Phi_{\mathbf{d}})]$ . Suppose  $\beta > 0$  is feasible. Recall that  $\rho$  is such that  $r = 2^{k-1} \ln 2 - \rho$ .

LEMMA 2. *W.h.p. the degree sequence  $\mathbf{d}$  chosen from  $\mathbf{D}$  is such that*

$$E[Z_{g,\beta}(\Phi_{\mathbf{d}})] \sim E[Z_\beta(\Phi_{\mathbf{d}})] \\ = \exp\left[\frac{n}{2^k} \left(2\rho - \ln 2 - (1 - \beta)\ln(1 - \beta) - \beta + O_k\left(\frac{k}{2^k}\right)\right)\right].$$

PROOF. Choose and fix a degree sequence  $\mathbf{d}$ . We need to compute the probability that some  $\sigma \in \{0,1\}^V$  is a good  $\beta$ -heavy solution. By symmetry, we may assume that  $\sigma = \mathbf{1}$  is the all-true assignment. Then  $\sigma$  is a solution iff every clause contains both a positive and a negative literal. Since the signs of the literals are chosen for all  $m$  clauses independently, we see that

$$\mathbb{P}[\sigma \in \mathcal{S}(\Phi_{\mathbf{d}})] = (1 - 2^{1-k})^m. \quad (16)$$

Given that  $\sigma$  is a solution, the number  $X$  of critical clauses has distribution  $\text{Bin}(m, k/(2^{k-1} - 1))$ , because whether a clause is critical depends on its signs only. As in the proof of Proposition 2, to determine the probability that  $\sigma$  is  $\beta$ -heavy we need to solve an occupancy problem:  $X$  balls representing the critical clauses are tossed randomly into  $n$  bins representing the variables. However, this time the bins have *capacities*: the bin representing  $x \in V$  can hold no more than  $\min\{3k, d_x\}$  balls in total. Thus, we need to compute the probability that under these constraints, exactly  $(1 - \beta)2^{-k}n$  bins are empty. This amounts to a rather non-trivial counting problem, but for a random degree sequence  $\mathbf{d}$  the probability differs from the formula obtained in Proposition 2 only by an error term that decays exponentially in  $k$ . More precisely,

$$\begin{aligned} & \mathbb{P}[\sigma \in \mathcal{S}_{\beta}(\Phi_{\mathbf{d}}) | \sigma \in \mathcal{S}(\Phi_{\mathbf{d}})] \\ &= \exp\left(-\frac{n}{2^k} \left[(1 - \beta)\ln(1 - \beta) - \beta - O_k(k/2^k)\right]\right). \end{aligned} \quad (17)$$

Let us provide some intuition why this is. The bin capacities are such that w.h.p. most bins can hold about  $kr = k2^{k-1} \ln 2 + O_k(k)$  balls. By comparison, the total number of balls is  $X \sim_k mk/(2^{k-1} - 1) \sim_k nk \ln 2$  w.h.p. In effect, the expected number of balls that a typical bin receives is about  $k \ln 2$ , way smaller than the capacity of that bin. Indeed, since the number of balls that are received by a typical bin is approximately  $\text{Bin}(kr, \frac{nk \ln 2}{km}) \approx \text{Bin}(kr, 2^{-k+1})$ , the number of balls can be approximated well by a  $\text{Po}(\lambda)$  distribution (with  $\lambda = kr/(2^{k-1} - 1) \sim_k k \ln 2$ ). Thus, the probability that a bin remains empty is close to  $\exp(-\lambda)$ , which was the probability of the same event in the experiment without capacities. The technical details of this argument are quite delicate, as the fluctuations of the capacities need to be controlled *very* carefully.

Finally, similar arguments as in the proof of Proposition 3 yield  $\mathbb{P}[\sigma \in \mathcal{S}_{g,\beta}(\Phi_{\mathbf{d}}) | \sigma \in \mathcal{S}_{\beta}(\Phi_{\mathbf{d}})] = 1 - o(1)$ . Thus, the assertion follows from (16)–(17).  $\square$

We now turn to the second moment. Fix some  $\sigma \in \{0,1\}^V$ , say  $\sigma = \mathbf{1}$ . Let  $Z_{g,\beta}(t, \sigma)$  denote the number of good  $\tau \in \mathcal{S}(\Phi_{\mathbf{d}})$  at distance  $t$  from  $\sigma$ . Using the linearity of expectation and recalling that the set of NAE-solutions is symmetric with respect to inversion, we obtain

$$\begin{aligned} & \mathbb{E}[Z_{g,\beta}(\Phi_{\mathbf{d}}) | \sigma \in \mathcal{S}_{g,\beta}(\Phi_{\mathbf{d}})] \\ & \leq 2 \sum_{0 \leq t \leq n/2} \mathbb{E}[Z_{g,\beta}(t, \sigma) | \sigma \in \mathcal{S}_{g,\beta}(\Phi_{\mathbf{d}})]. \end{aligned} \quad (18)$$

Let  $I = \left\{t \in \mathbf{Z} : (\frac{1}{2} - 2^{-k/3})n \leq t \leq n/2\right\}$ . The first two conditions from Definition 1 ensure that given that  $\sigma$  is good, so with certainty we have

$$\sum_{t \leq 0.01n} Z_{g,\beta}(t, \sigma) \leq |\mathcal{C}(\sigma)| \leq \mathbb{E}[Z_{\beta}(\Phi_{\mathbf{d}})]$$

and

$$\sum_{0.01n < t < (\frac{1}{2} - 2^{-k/3})n} Z_{g,\beta}(t, \sigma) = 0.$$

Hence, Lemma 2 and (18) yield

$$\begin{aligned} & \mathbb{E}[Z_{g,\beta}(\Phi_{\mathbf{d}}) | \sigma \in \mathcal{S}_{g,\beta}(\Phi_{\mathbf{d}})] \\ & \leq (2 + o(1)) \mathbb{E}[Z_{g,\beta}(\Phi_{\mathbf{d}})] + 2 \sum_{t \in I} \mathbb{E}[Z_{g,\beta}(t, \sigma) | \sigma \in \mathcal{S}_{g,\beta}(\Phi_{\mathbf{d}})]. \end{aligned} \quad (19)$$

This reduces the proof to the analysis of the “central terms” with  $t \in I$ . The result of this is

LEMMA 3. *There is a constant  $C' = C'(k) \geq 1$  such that for a random  $\mathbf{d}$  we have w.h.p.*

$$\sum_{t \in I} \mathbb{E}[Z_{g,\beta}(t, \sigma) | \sigma \in \mathcal{S}_{g,\beta}(\Phi_{\mathbf{d}})] \leq C' \cdot \mathbb{E}[Z_{g,\beta}(\Phi_{\mathbf{d}})]. \quad (20)$$

PROOF SKETCH. This is technically the most challenging bit of this work. The argument boils down to estimating the probability that two random  $\sigma, \tau \in \{0,1\}^n$  with  $\text{dist}(\sigma, \tau)/n = \alpha \in [\frac{1}{2} - 2^{-k/3}, \frac{1}{2}]$  simultaneously are good  $\beta$ -heavy solutions. To compute this probability, we need to analyze the interplay of two occupancy problems as in the proof of Lemma 2 with respect to the same degree sequence  $\mathbf{d}$ .

More precisely, let  $B = \bigcup_{x \in V} \{x\} \times \{1, \dots, d_x\}$  be a set of  $km$  “balls”. Generating  $\Phi_{\mathbf{d}}$  is equivalent to drawing a random bijection  $\pi : [m] \times [k] \rightarrow B$ , with  $\pi(i, j) = (x, l)$  indicating that  $x$  is the underlying variable of the  $j$ th literal of clause  $i$ , and independently choosing a map  $\mathbf{s} : [m] \times [k] \rightarrow \{\pm 1\}$  indicating the signs. Further, we represent the occupancy problems for  $\sigma, \tau$  by two “colorings”  $g_{\sigma}, g_{\tau} : B \rightarrow \{\text{red}, \text{blue}\}$ , with  $g_{\sigma}(x, l) = \text{red}$  indicating that the  $l$ th position in bin  $x$  is occupied under  $\sigma$  (and analogously for  $\tau$ ). We compute the probability  $p(\alpha, g_{\sigma}, g_{\tau})$  that  $\pi, \mathbf{s}$  induce a formula in which

- literal  $(i, j)$  supports clause  $i$  under  $\sigma$  iff  $g_{\sigma} \circ \pi(i, j) = \text{red}$ , and similarly for  $\tau$ .
- both  $\sigma, \tau$  are good  $\beta$ -heavy solutions.

The result is that for any  $g_{\sigma}, g_{\tau}$  the “success probability” is *minimized* at  $\alpha = 1/2$ . Quantitatively,

$$\frac{p(\alpha, g_{\sigma}, g_{\tau})}{p(1/2, g_{\sigma}, g_{\tau})} = \exp\left[O_k(k^4/2^k)(\alpha - 1/2)^2 n\right] \quad \text{for any } g_{\sigma}, g_{\tau}. \quad (21)$$

On the other hand, the total *number* of assignment pairs satisfies

$$\begin{aligned} & \frac{|\{(\sigma, \tau) : \text{dist}(\sigma, \tau) = \alpha n\}|}{|\{(\sigma, \tau) : \text{dist}(\sigma, \tau) = n/2\}|} \\ &= \binom{n}{\alpha n} / \binom{n}{n/2} \\ &= \exp(-(4 - o_k(1))(\alpha - 1/2)^2 n), \end{aligned} \quad (22)$$

which is *maximized* at  $\alpha = 1/2$ . Combining (21) and (22), we see that for any two colorings  $g_{\sigma}, g_{\tau}$  the dominant contribution to the second moment stems from  $\alpha = \frac{1}{2} + O(1/\sqrt{n})$ , i.e., from “perfectly decorrelated”  $\sigma, \tau$ . The assertion follows by evaluating the contribution of such  $\alpha$  explicitly and summing over  $g_{\sigma}, g_{\tau}$ .  $\square$

**Acknowledgment.** The first author thanks Dimitris Achlioptas and Lenka Zdeborová for helpful discussions on the second moment method and the statistical mechanics work on random CSPs.

## 6. REFERENCES

- [1] D. Achlioptas, A. Coja-Oghlan: Algorithmic barriers from phase transitions. Proc. 49th FOCS (2008) 793–802.
- [2] D. Achlioptas, J.H. Kim, M. Krivelevich, P. Tetali: Two-coloring random hypergraphs. Random Structures and Algorithms **18** (2002), 249–259.
- [3] D. Achlioptas, C. Moore: Random  $k$ -SAT: two moments suffice to cross a sharp threshold. SIAM Journal on Computing **36** (2006) 740–762.
- [4] D. Achlioptas, A. Naor: The two possible values of the chromatic number of a random graph. Annals of Mathematics **162** (2005) 1333–1349.
- [5] D. Achlioptas, A. Naor, Y. Peres: Rigorous location of phase transitions in hard optimization problems. Nature **435** (2005) 759–764.
- [6] 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.
- [7] D. Achlioptas, F. Ricci-Tersenghi: On the solution space geometry of random constraint satisfaction problems. Proc. 38th STOC (2006) 130–139.
- [8] V. Chvátal, B. Reed: Mick gets some (the odds are on his side). Proc. 33th FOCS (1992) 620–627.
- [9] A. Coja-Oghlan: A better algorithm for random  $k$ -SAT. SIAM J. Computing **39** (2010) 2823–2864.
- [10] A. Coja-Oghlan, C. Efthymiou: On independent sets in random graphs. Proc. 22nd SODA (2011) 136–144.
- [11] A. Coja-Oghlan, A. Frieze: Random  $k$ -SAT: the limiting probability for satisfiability for moderately growing  $k$ . Electronic Journal of Combinatorics **15** (2008) N2.
- [12] A. Coja-Oghlan, L. Zdeborová: The condensation transition in random hypergraph 2-coloring. Proc. 23rd SODA (2012), to appear.
- [13] H. Connamacher, M. Molloy: The exact satisfiability threshold for a potentially intractable random constraint satisfaction problem. Proc. 45th FOCS (2004) 590–599.
- [14] L. Dall’Asta, A. Ramezanpour, R. Zecchina: Entropy landscape and non-Gibbs solutions in constraint satisfaction problems. Phys. Rev. E **77**, 031118 (2008).
- [15] H. Daudé, M. Mézard, T. Mora, R. Zecchina: Pairs of SAT-assignments in random Boolean formulae. Theoretical Computer Science **393** (2008) 260–279.
- [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] C. Efthymiou: A simple algorithm for random colouring  $G(n, d/n)$  using  $(2 + \epsilon)d$  colours. Proc. 23rd SODA (2012), to appear.
- [19] U. Feige, E. Mossel, D. Vilenchik: Complete convergence of message passing algorithms for some satisfiability problems. Proc. 10th RANDOM (2006) 339–350.
- [20] P. Flajolet, R. Sedgewick: Analytic Combinatorics. Cambridge University Press, Cambridge, 2009.
- [21] S. Franz, M. Leone: Replica bounds for optimization problems and diluted spin systems. J. Statist. Phys. **111** (2003) 535–564.
- [22] E. Friedgut: Hunting for sharp thresholds. Random Struct. Algorithms **26** (2005) 37–51.
- [23] A. Frieze, N. Wormald: Random  $k$ -Sat: a tight threshold for moderately growing  $k$ . Combinatorica **25** (2005) 297–305.
- [24] A. Gerschenfeld, A. Montanari: Reconstruction for models on random graphs. Proc. 48th FOCS (2007) 194–204.
- [25] A. Goerdts: A threshold for unsatisfiability. Proc. 17th MFCS (1992) 264–274.
- [26] L. Kirousis, E. Kranakis, D. Krizanc, Y. Stamatou: Approximating the unsatisfiability threshold of random formulas. Random Structures Algorithms **12** (1998) 253–269.
- [27] 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.
- [28] S. Mertens, M. Mézard, R. Zecchina: Threshold values of random  $K$ -SAT from the cavity method. Random Struct. Alg. **28** (2006) 340–373.
- [29] M. Mézard, G. Parisi, R. Zecchina: Analytic and algorithmic solution of random satisfiability problems. Science **297** (2002) 812–815.
- [30] A. Montanari, R. Restrepo, P. Tetali: Reconstruction and clustering in random constraint satisfaction problems. SIAM J. Discrete Math. **25** (2011) 771–808.
- [31] D. Panchenko, M. Talagrand: Bounds for diluted mean-fields spin glass models. Probab. Theory Relat. Fields **130** (2004) 319–336.
- [32] B. Pittel, G. Sorkin: The satisfiability threshold for  $k$ -XORSAT. Preprint (2011).