Spines of random Constraint Satisfaction Problems: definition and impact on Computational Complexity

Stefan Boettcher*, Gabriel Istrate and Allon G. Percus†

1 Introduction

The major promise of *phase transitions in combinatorial problems* was to shed light on the "practical" algorithmic complexity of combinatorial problems. A possible connection has been highlighted by the results (based on experimental evidence and nonrigorous arguments from statistical mechanics) of Monasson et al. [1, 2]. Studying a version of random satisfiability that "interpolates" between 2-SAT and 3-SAT, they concluded that the order of the phase transition, combinatorially expressed by continuity of an order parameter called the *backbone*, might have algorithmic implications for the complexity of the important class of *Davis-Putnam-Longman-Loveland (DPLL) algorithms* [3]. A discontinuous or first-order transition was symptomatic of exponential complexity, whereas a continuous or second-order transition correlated with polynomial complexity.

It is well understood by now that this connection is limited. For instance, k-XOR-SAT is a problem believed, based on arguments from statistical mechanics [4], to have a first-order phase transition. But it is easily solved by a polynomial algorithm, Gaussian elimination.

One way to clarify the connection between phase transitions and computational complexity is to formalize the underlying intuition in a purely combinatorial way, devoid of any physics considerations. First-order phase transitions amount to a discontinuity in the (normalized) size of the backbone. For random k-SAT [5], and more specifically for the optimization problem MAX-k-SAT, the backbone has a combinatorial interpretation: it is the the set of literals that are "frozen" (assume the same value) in all *optimal* assignments. Intuitively, a large backbone size has implications for the complexity of finding such assignments: all literals in the backbone require well-defined values in order to satisfy the formula, but an algorithm assigning variables in an iterative fashion has very few ways to know what the "right" values to assign are. In the case in a first-order phase transition, the backbone of formulas just above the transition contains with high probability a fraction of the literals that is bounded away from zero. DPLL algorithms would then misassign a variable having $\Omega(n)$ height in the tree representing the behavior of the algorithm, forcing it to backtrack on the given variable. Assuming the algorithm cannot significantly "reduce" the size of the explored portion of this tree, a first-order phase transition would then w.h.p imply a $2^{\Omega(n)}$ lower bound for the running time of DPLL on random instances located slightly above the transition.

There exists, however, a significant flaw in the heuristic argument above: the backbone is defined with respect to *optimal* solutions, and would seem to imply that it is difficult to find solutions to, e.g., MAX-K-SAT using algorithms that assign variables iteratively. But why should the continuity/discontinuity of the backbone be the relevant predictor for the complexity of the (often easier) *decision* problem, which is what DPLL algorithms try to solve anyway?

*stb@physics.emory.edu Emory University, Atlanta, GA 30322

†{istrate,percus}@lanl.gov Los Alamos National Laboratory, Los Alamos, NM 87545. Correspondence to: Gabriel Istrate.

Fortunately, it turns out that the intuition of the previous argument also holds for a different order parameter, a "weaker" version of the backbone called the *spine*, introduced in [6] in order to prove that random 2-SAT has a second-order phase transition. Unlike the backbone, the spine is defined in terms of the *decision* problem, hence it could conceivably have a larger impact on the complexity of these problems. Of course, the same caveat applies as for the backbone: we are referring to complexity with respect to classes of algorithms weaker than polynomial time computations, and that in particular are not strong enough to capture the polynomial time Gaussian elimination algorithm for k-XOR-SAT.

We aim in this paper to provide evidence that the behavior of the spine, rather than the backbone, impacts the complexity of the underlying decision problem. To accomplish this:

1. We discuss the proper definition of the backbone/spine for random CSP.

2. We formally establish a simple connection between a discontinuity in the (relative size of the) spine at the threshold and the resolution complexity of random satisfiability problems. In a nutshell, a necessary and sufficient condition for the existence of a discontinuity in the spine is the existence of a $\Omega(n)$ lower bound (w.h.p.) on the size of minimally unsatisfiable subformulas of a random (unsatisfiable) subformula. But standard methods from proof complexity [7] imply that (in conjunction with the expansion of the formula hypergraph, *independent* of the precise definition of the problem at hand) in all cases where we can prove the existence of a first-order phase transition, such problems have a $2^{\Omega(n)}$ lower bound on their resolution complexity (and hence the complexity of DPLL algorithms as well [3]). In contrast, we show (Theorem 1) that, for *any generalized satisfiability problem*, a second-order phase transition implies, for every $\alpha > 0$, a $O(2^{\alpha \cdot n})$ upper bound on the resolution complexity of their random instances (in the region where most formulas are unsatisfiable)

3. We give a sufficient condition (Theorem 2) for the existence of a discontinuous jump in the size of the spine. We then show (Theorem 3) that all problems whose constraints have no implicates of size at most two satisfy this condition. Qualitatively, our results suggest that all satisfiability problems with a second-order phase transition in the spine are "like 2-SAT".

4. Finally, we present some experimental results that attempt to clarify the issue whether the backbone and the spine can behave differently at the phase transition. The *Graph bipartition problem* is one case where this seems to happen. In contrast, the backbone and spine of *random 3-coloring* seem to have similar behavior.

A note on the significance of our results: a discontinuity in the spine is weaker than a first-order phase transition (i.e., a discontinuity in the size of the backbone). Also, unlike the backbone, the spine has no physical interpretation. But this is *not* our intention: we have seen that the argument connecting the backbone size and the complexity of decision problems is problematic. What we rigorously show (with no physics considerations in mind) is that *the intuitive argument holds for the spine order parameter*. Moreover, the last section of the paper presents experimental work suggesting that *the backbone and the spine can behave differently*.

2 Preliminaries

Throughout the paper we will assume familiarity with the general concepts of phase transitions in combinatorial problems (see e.g., [8]), random structures [9], proof complexity [10]. Some papers whose concepts and methods we use in detail (and we assume greater familiarity with) include [11], [12], [7]. We will use the model of random constraint satisfaction from Molloy [13]:

Definition 1 *Let $\mathcal{D} = \{0, 1, \ldots, t-1\}$, $t \geq 2$ be a fixed set. Consider all $2^{t^k} - 1$ possible nonempty sets of constraints (relations) on k variables $X_1, \ldots, X_k$ with values taken from $\mathcal{D}$. Let $\mathcal{C}$ be such a nonempty set of constraints.*

A random formula from CSP$_{n,m}(\mathcal{C})$ is specified by the following procedure:

- *n is the number of variables.*
- *m is the number of clauses, chosen by the following procedure: first select, uniformly at random and with replacement, m hyperedges of the complete k-uniform hypergraph on n variables.*
- *for each hyperedge, choose a random ordering of the variables involved. Choose a random constraint from $\mathcal{C}$ and apply it on the list of (ordered) variables.*

We use the notation SAT($\mathcal{C}$) (instead of CSP($\mathcal{C}$)) when t=2. Also, for Φ an instance of CSP($\mathcal{C}$) we denote by $opt_{\mathcal{C}}(\Phi)$ the smallest number of constraints left unsatisfied by some assignment.

Just as in random graphs [9], under fairly liberal conditions one can use the *constant probability model* instead of the *counting model* from the previous definition. The interesting range of the parameter m is when the ratio m/n is a constant, c, the *constraint density* (details are left for the final version of the paper). The original investigation of the order of the phase transition in k-SAT used an order parameter called *the backbone*,

$$B(\Phi) = \{x \in Lit \mid \exists\, \lambda \in \{0,1\} : \forall\, \Xi \in \mathrm{MAXSAT}(\Phi), \Xi(x) = \lambda\}, \tag{1}$$

or more precisely the backbone fraction f, the fraction of the n variables that belong to $B(\Phi)$.

Bollobás et al. [6] have investigated the order of the phase transition in k-SAT (for $k = 2$) under a different order parameter, a "monotonic version" of the backbone called *the spine*.

$$S(\Phi) = \{x \in Lit \mid \exists\, \Xi \subseteq \Phi : \Xi \in k\mathrm{-SAT}, \Xi \wedge \{\overline{x}\} \in \overline{k\mathrm{-SAT}}\}. \tag{2}$$

They showed that random 2-SAT has a continuous (second-order) phase transition: the size of the spine, normalized by dividing it by the number of variables, approaches zero (as $n \to \infty$) for $c < c_{2-\mathrm{SAT}} = 1$, and is continuous at $c = c_{2-\mathrm{SAT}}$. By contrast, nonrigorous arguments from statistical mechanics [5] imply that for 3-SAT the backbone jumps discontinuously from zero to positive values at the transition point $c = c_{3-\mathrm{SAT}}$ (a first-order phase transition).

3 How to define the backbone/spine for random CSP (and beyond)

We would like to extend the concepts of backbone and spine to general constraint satisfaction problems. Certain differences between the case of random k-SAT and the general case force us to employ an alternative definition of the backbone/spine. The most obvious is that formula (2) involves negations of variables, unlike Molloy's model. Also, these definitions are inadequate for problems whose solution space presents a relabelling symmetry, such as the case of *graph coloring* where the set of (optimal) colorings is closed under permutations of the colors. Due to this symmetry, no variable can be frozen in this way.

The new definitions have to retain as many of the properties of the backbone/spine as possible. In particular, the new definitions must give rise to *order parameters*, i.e., quantities that are zero up to the critical value of the control parameter (in our case constraint density c) and positive above it. The formal statement that we wish to extend to CSP($\mathcal{C}$) is presented next for the spine:

Lemma 3.1 *Let c be an arbitrary* constant *value for the constraint density function.*

1. *If $c < \underline{lim}_{n\to\infty} c_{k-SAT}(n)$ then $\lim_{n\to\infty} \frac{|S(\Phi)|}{n} = 0$.*
2. *If for some c there exists $\delta > 0$ such that w.h.p. (as $n \to \infty$) $\frac{|S(\Phi)|}{n} > \delta$ then $\lim_{n\to\infty} \mathrm{Prob}[\Phi \in SAT] = 0$, that is $c > \overline{lim}_{n\to\infty} c_{k-SAT}(n)$.*

Our solution is to define the backbone/spine of a random instance of CSP($\mathcal{C}$) slightly differently.

Definition 2

$$B(\Phi) = \{x \in Var \mid \exists\, C \in \mathcal{C} : x \in C, opt_{\mathcal{C}}(\Phi \cup C) > opt_{\mathcal{C}}(\Phi)\},$$
$$S(\Phi) = \{x \in Var \mid \exists\, C \in \mathcal{C} \text{ and } \Xi \subseteq \Phi : x \in C, \Xi \in \mathrm{CSP}, \Xi \wedge C \in \overline{\mathrm{CSP}}\}.$$

For k-CNF formulas whose (original) backbone/spine contains at least three literals, a variable x is in the (new version of the) backbone/spine if and only if either x or $\overline{x}$ were present in the old version. In particular the new definition does not change the order of the phase transition of random k-SAT.

Alternatively, in studying 3-colorability of random graphs $G = (V, E)$, Culberson and Gent [14] define

$$S(G) = \{(x, y) \in V^2 \mid \exists\, H \subseteq G : H \in \text{3-COL}, H \cup E(x, y) \in \overline{\text{3-COL}}\},$$

so one may define the relative backbone and spine sizes in terms of *constraints* rather than *variables*.

Definition 3

$$B(\Phi) = \{C \in \mathcal{C} \mid opt_{\mathcal{C}}(\Phi \cup C) > opt_{\mathcal{C}}(\Phi)\},$$
$$S(\Phi) = \{C \in \mathcal{C} \mid \exists\, \Xi \subseteq \Phi : \Xi \in \mathrm{CSP}, \Xi \wedge C \in \overline{\mathrm{CSP}}\}.$$

Since we are looking at a combinatorial definition, with no physics considerations in mind, the only principled way to choose between the two types of order parameters (one based on variables, the other based on constraints) is to look at the class of algorithms we are concerned with. In the case of random constraint satisfaction problems (and DPLL algorithms) it is variables that get assigned values, so Definition 2 is preferred. On the other hand, we will see an example in a later section (the case of *graph bipartitioning*) where it makes more sense to use Definition 3.

Finally, note that if the backbone/spine of an instance of CSP($\mathcal{C}$) (in the sense of definition 2) has size u, then the backbone/spine (in the sense of definition 3) has size $O(u^k)$. It follows readily that the discontinuity of the second version of backbone/spine implies the discontinuity of the corresponding first version. This will notably be the case in our experimental study of 3-COL.

4 Spine discontinuity and resolution complexity of random CSP

Definition 4 *Let $\mathcal{C}$ be such that SAT($\mathcal{C}$) has a sharp threshold. Problem SAT($\mathcal{C}$) has a* discontinuous spine *if there exists $\eta > 0$ such that for every sequence $m = m(n)$ we have*

$$\lim_{n\to\infty} \mathop{\mathrm{Prob}}_{m=m(n)}[\Phi \in SAT] = 0 \Rightarrow \lim_{n\to\infty} \mathop{\mathrm{Prob}}_{m=m(n)}[\frac{|S(\Phi)|}{n} \geq \eta] = 1. \quad (3)$$

If, on the other hand, for every $\epsilon > 0$ there exists $m^{\epsilon}(n)$ with

$$\lim_{n\to\infty} \mathop{\mathrm{Prob}}_{m=m^{\epsilon}(n)}[\Phi \in SAT] = 0 \text{ and } \lim_{n\to\infty} \mathop{\mathrm{Prob}}_{m=m^{\epsilon}(n)}[\frac{|S(\Phi)|}{n} \geq \epsilon] = 0 \quad (4)$$

we say that SAT($\mathcal{C}$) has a continuous spine.

Claim 1 *Let Φ be minimally unsatisfiable, and let x be a literal that appears in Φ. Then $x \in S(\Phi)$.*

Corollary 1 *k-SAT, $k \geq 3$ has a discontinuous spine.*

The *resolution complexity* of an instance Φ of SAT($\mathcal{C}$) is defined as the resolution complexity of the formula obtained by converting each constraint of Φ to CNF-form. A simple observation is that a continuous spine has implications for resolution complexity:

Theorem 1 *Let $\mathcal{C}$ be a set of constraints such that SAT($\mathcal{C}$) has a continuous spine. Then for every value of the constraint density $c > \overline{lim}_{n\to\infty} c_{SAT(\mathcal{C})}(n)$, and* every *$\alpha > 0$, random formulas of constraint density c have w.h.p. resolution complexity $O(2^{k\cdot\alpha\cdot n})$.*

Definition 5 *For a formula F define $c^*(F) = \max\{\frac{|Constraints(G)|}{|Var(G)|} : \emptyset \neq G \subseteq F\}$.*

The next result gives a sufficient condition for a generalized satisfiability problem to have a discontinuous spine. Interestingly, it is one condition studied in [13].

Theorem 2 *Let $\mathcal{C}$ be such that SAT($\mathcal{C}$) has a sharp threshold. If there exists $\epsilon > 0$ such that for every minimally unsatisfiable formulas F it holds that $c^*(F) > \frac{1+\epsilon}{k-1}$, then SAT($\mathcal{C}$) has a discontinuous spine.*

One can give an explicitly defined class of satisfiability problems for which the previous result applies:

Theorem 3 *Let $\mathcal{C}$ be such that SAT($\mathcal{C}$) has a sharp threshold. If* no *clause $C \in \mathcal{C}$ has an implicate of length at most 2 then*

1. *for every minimally unsatisfiable formula F, $c^*(F) \geq \frac{2}{2k-3}$. Therefore SAT($\mathcal{C}$) satisfies the conditions of the previous theorem, i.e., it has a discontinuous spine.*

2. *Moreover SAT($\mathcal{C}$) also has $2^{\Omega(n)}$ resolution complexity*[1].

The condition in the theorem is violated (as expected) by random 2-SAT, as well as by the random version of the NP-complete problem 1-in-k SAT. This problem can be represented as CSP($\mathcal{C}$), for $\mathcal{C}$ a set of 2^k constraints (corresponding to all ways to negate some of the variables) and has a rigorously determined "2-SAT like" location of the transition point [17]. But the formula $C(x_1, x_2, \ldots, x_{k-1}, x_k) \wedge C(\overline{x_k}, x_{k+1}, \ldots, x_{2k-2}, x_1) \wedge C(\overline{x_1}, x_{2k-1}, \ldots, x_{3k-3}, \overline{x_k}) \wedge C(x_k, x_{3k-2}, \ldots, x_{4k-4}, x_1)$, where C is the constraint "1-in-k", is minimally unsatisfiable but has clause/variable ratio $1/(k-1)$ and implicates $\overline{x_1} \vee \overline{x_k}$ and $x_1 \vee x_k$.

4.1 Threshold location and discontinuous spines

Molloy [13] has shown (Theorem 3 in his paper) that the condition that turned out to be sufficient for the existence of a phase transition has implications for the location of the threshold. A natural question therefore arises: is it possible to read the order of the phase transition from the location of the threshold?

We cannot answer this question in full. However, we have already seen two problems that do not satisfy our sufficient condition for a discontinuous spine: random 2-SAT, for which the transition has been proven to be of second order [6], and random 1-in-k SAT, for which we believe a similar result holds [17]. Molloy's result does not provide the correct location of the threshold for these two problems. It is, however, striking, that for both problems the actual location of the threshold is *twice* the value given by Theorem 3 in [13], at clause/variable ratio $2/k(k-1)$. We give, therefore the following observation, a variant of Molloy's result (proved in the journal version of the paper).

[1]this result subsumes some of the results in [15]. Related results have been given independently in [16]

Proposition 1 *Modify the random model from Definition 1 to:*

- *allow application of constraints to both variables and negated variables.*
- *only allow constraints that cannot be made equal via negation of some of their variables.*

Denote by $CSP_{neg}(\mathcal{C})$ the new random model.

Let $\mathcal{C}$ be such that for every minimally unsatisfiable subformula F whose constraints are drawn from F the ratio of constraints to variables of F is at least $\frac{1+\epsilon}{k-1}$, for some $\epsilon > 0$. Then there is a constant $\delta > 0$ such that $CSP_{neg}(\mathcal{C})$ is a.s satisfiable for $m \leq \frac{2}{k(k-1)} \cdot (1+\delta)$.

5 Beyond random satisfiability: comparing the behavior of the backbone and spine

In this section we investigate empirically the continuity of the backbone for two graph problems, random three coloring (3-COL) and graph bipartition (GBP).

We consider a large number of instances of random graphs, of sizes up to $n = 1024$ and over a range of mean degree values near the threshold. for each instance we determine the backbone fraction f.

Culberson and Gent [14] have shown experimentally that the 3-COL spine (as defined in Definition 3) exhibits a discontinuous transition. To be consistent with this study (and since it leads to a stronger result anyhow) we use the backbone from the same definition. We employ a rapid heuristic called *extremal optimization* [18] that, based on testbed comparisons with an exact algorithm, yields an excellent approximation of f around the critical region. Figure 1 shows f as a function of mean degree. Above the threshold, for 3-COL (Fig. 1a) f does not appear to vanish, suggesting a discontinuous large n backbone. Culberson and Gent have speculated that at the 3-COL threshold, although their spine is discontinuous, the backbone might be *continuous*. Our numerical results suggest otherwise.

We next study graph bipartitioning (GBP). This problem cannot, strictly speaking, be cast in the setup of random constraint satisfaction problems from Definition 1, since not every partition of vertices of G is allowed. It can be cast to a variant of this model (with variables associated to nodes, values associated to each partition and constraint "$x = y$" associated to the edge between the corresponding vertices) but we must add the additional requirement that all satisfying assignments contain an equal number of ones and zeros. Thus the complexity-theoretic observations of Section 4 do not automatically apply to it. We can, however, give a "DPLL-like" class of algorithms for GBP, that assigns vertices (variables) in *pairs*, one to each partition. This class of algorithms motivates investigating the backbone/spine under the model in Definition 3. It is easy to see that the spine of a GBP instance contains all edges belonging to a connected component of size larger than $n/2$. Since the phase transition in GBP takes place where the giant component becomes larger than $n/2$, GBP has a discontinuous spine. The backbone (Fig. 1b), on the other hand, appears to remains continuous, vanishing at large n on both sides of the threshold.

One obvious question raised by the previous result is whether the discontinuity of the spine in GBP really has computational implications for the complexity of deciding whether a perfect bipartition exists. After all, unlike 3-COL, GBP can easily be solved in polynomial time by dynamic programming. This situation, however, is similar to that of XOR-SAT, where a polynomial algorithm exists but the complexity of *resolution proofs/DPLL algorithms* is exponential.

The class of "DPLL-like" algorithms outlined for GBP can no longer be simulated in the straightforward manner by *resolution proofs*, however it can be simulated using proof systems $Res(k)$ that are extensions of resolution [19]. Some of the hardness results for resolution extend to these more powerful

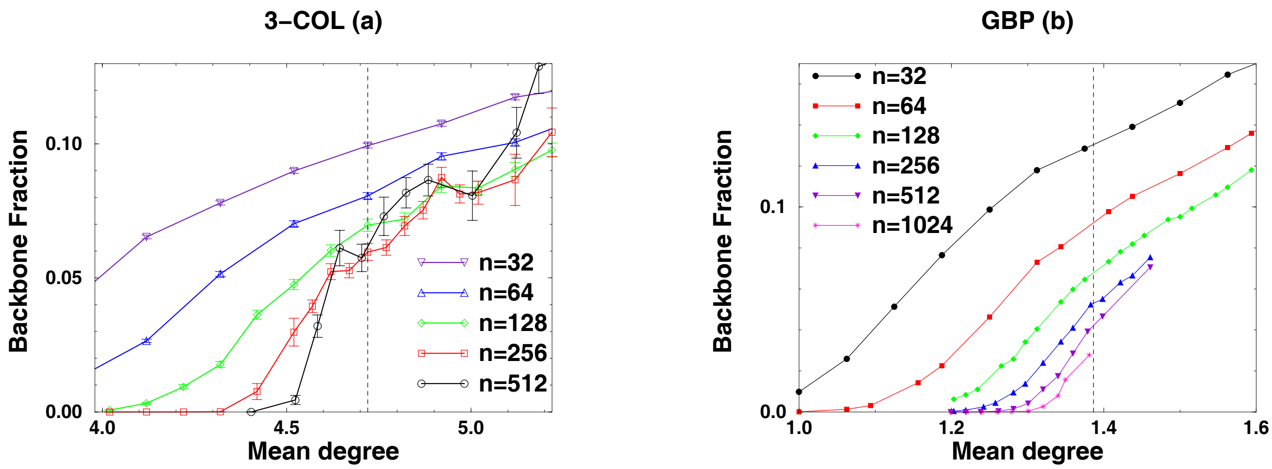

Figure 1: Plot of the estimated backbone fraction (a) for 3-COL and (b) for the GBP, on random graphs, as a function of mean degree α. For 3-COL, the systematic error based on benchmark comparisons with random graphs is negligible compared to the statistical error bars; for the GBP, f is found by exact enumeration. The thresholds $\alpha \approx 4.73$ for 3-COL and $\alpha = 2\ln 2$ for the GBP are shown by dashed lines.

proof systems, and in [20] we investigate the extent to which the results of this paper apply to this class of proof systems. These preliminary results imply that, indeed, the discontinuity of the spine *does* have computational implications for GBP.

6 Discussion

We have shown that the existence of a discontinuous spine in a random satisfiability problem is often correlated with a $2^{\Omega(n)}$ peak in the complexity of resolution/DPLL algorithms at the transition point. The underlying reason is that the two phenomena (the jump in the order parameter and the resolution complexity lower bound) have common causes.

The example of random k-XOR-SAT shows that a general connection between first-order phase transition and the complexity of the decision problems is hopeless: Ricci-Tersenghi et al. [4] have presented a non-rigorous argument using the replica method that shows that this problem has a first-order phase transition, and we can formally show (as a direct consequence of Theorem 3) the following weaker result:

Proposition 2 *Random k-XOR-SAT, $k \geq 3$, has a discontinuous spine.*

However, experimental evidence in previous section suggests that the backbone and the spine do not always behave in the same way. Therefore our results (and the work in progress mentioned above) suggest that the continuity/discontinuity of the spine, rather than the backbone, is a predictor for the complexity of the *restricted classes* of decision algorithms that can be simulated by “resolution-like” proof systems.

References

[1] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Determining computational complexity from characteristic phase transitions. *Nature*, 400(8):133–137, 1999.

[2] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. $(2+p)$-SAT: Relation of typical-case complexity to the nature of the phase transition. *Random Structures and Algorithms*, 15(3–4):414–435, 1999.

[3] P. Beame, R. Karp, T. Pitassi, and M. Saks. The efficiency of resolution and Davis-Putnam procedures. *SIAM Journal of Computing*, 31(4):1048–1075, 2002.

[4] F. Ricci-Tersenghi, M. Weight, and R. Zecchina. Simplest random k-satisfiability problem. *Physical Reviews E*, 63:026702, 2001.

[5] R. Monasson and R. Zecchina. Statistical mechanics of the random k-SAT model. *Physical Review E*, 56:1357, 1997.

[6] B. Bollobás, C. Borgs, J.T. Chayes, J. H. Kim, and D. B. Wilson. The scaling window of the 2-SAT transition. *Random Structures and Algorithms*, 18(3):201–256, 2001.

[7] E. Ben-Sasson and A. Wigderson. Short Proofs are Narrow:Resolution made Simple. *Journal of the ACM*, 48(2), 2001.

[8] O. Martin, R. Monasson, and R. Zecchina. Statistical mechanics methods and phase transitions in combinatorial optimization problems. *Theoretical Computer Science*, 265(1-2):3–67, 2001.

[9] B. Bollobás. *Random Graphs*. Academic Press, 1985.

[10] P. Beame and T. Pitassi. Propositional proof complexity: Past present and future. In *Current Trends in Theoretical Computer Science*, pages 42–70. 2001.

[11] E. Friedgut. Necessary and sufficient conditions for sharp thresholds of graph properties, and the k-SAT problem. with an appendix by J. Bourgain. *Journal of the A.M.S.*, 12:1017–1054, 1999.

[12] V. Chvátal and E. Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, 1988.

[13] M. Molloy. Models for random constraint satisfaction problems. In *Proceedings of the 32nd ACM Symposium on Theory of Computing*, 2002.

[14] J. Culberson and I. Gent. Frozen development in graph coloring. *Theoretical Computer Science*, 265(1-2):227–264, 2001.

[15] D. Mitchell. Resolution complexity of Random Constraints. In *Eigth International Conference on Principles and Practice of Constraint Programming*, 2002.

[16] M. Molloy and M. Salavatipour. The resolution complexity of random constraint satisfaction problems. In *Proceedings of the 44th Symposium on Foundations in Computer Science*, 2003.

[17] D. Achlioptas, A. Chtcherba, G. Istrate, and C. Moore. The phase transition in random 1-in-k SAT and NAE 3-SAT. In *Proceedings of the 13th ACM-SIAM Symposium on Discrete Algorithms*, 2001.

[18] S. Boettcher and A. Percus. Nature's way of optimizing. *Artificial Intelligence*, 119:275–286, 2000.

[19] J. Krajicek. On the weak pigeonhole principle. *Fundamenta Matematicae*, 170(1–3):123–140, 2001.

[20] G. Istrate. Descriptive complexity and first-order phase transitions. (in progress).

Appendix

6.1 Proof of Claim 1

Let C be a clause that contains x. By the minimal unsatisfiability of Φ, $\Phi \setminus \{C\} \in SAT$. On the other hand $\Phi \setminus \{C\} \wedge \{x\} \in \overline{SAT}$, otherwise Φ would also be satisfiable. Thus $x \in S(\Phi \setminus \{C\})$.

6.2 Proof of Corollary 1

To show that 3-SAT has a discontinuous spine it is enough to show that a random unsatisfiable formula contains w.h.p. a minimally unsatisfiable subformula containing a linear number of literals. A way to accomplish this is by using the two ingredients of the Chvátal-Szemerédi proof [12] that random 3-SAT has exponential resolution size w.h.p.

6.3 Proof of Theorem 1

Proof.

By the analog of Claim 1 for the general case, if $SAT(\mathcal{C})$ has a continuous spine) then for every $c > c_{SAT(\mathcal{C})}$ and for every $\alpha > 0$, minimally unsatisfiable subformulas of a random formula Φ with constraint density c have w.h.p. size at most $\alpha \cdot n$. Consider the backtrack tree of the natural DPLL algorithm (that tries to satisfies clauses one at a time) on such a minimally unsatisfiable subformula F. By the usual correspondence between DPLL trees and resolution complexity (e.g., [3], pp. 1) it yields a resolution proof of the unsatisfiability of Φ having size at most $2^{k \cdot \alpha \cdot n+1}$.

□

6.4 Proof of Theorem 2

Proof.

We first recall the following concept from [12]:

Definition 6 *Let $x, y > 0$. A k-uniform hypergraph with n vertices is* (x,y)-sparse *if every set of $s \leq xn$ vertices contains at most ys edges.*

We also recall Lemma 1 from the same paper.

Lemma 6.1 *Let $k, c > 0$ and y be such that $(k-1)y > 1$. Then w.h.p. the random k-uniform hypergraph with n vertices and cn edges is (x, y)-sparse, where $\epsilon = y - 1/(k-1)$, $x = (\frac{1}{2e}(\frac{y}{ce})^y)^{1/\epsilon}$.*

The critical observation is that the existence of a minimally unsatisfiable formula of size xn and with $c^*(F) > \frac{1+\epsilon}{k-1}$ implies that the k-uniform hypergraph associated to the given formula is *not* (x, y)-sparse, for $y = \frac{\epsilon}{k-1}$. But, according to Lemma 6.1, w.h.p. a random k-uniform hypergraph with cn edges is (x_0, y) sparse, for $x_0 = (\frac{1}{2e}(\frac{y}{ce})^y)^{1/\epsilon}$ (a direct application of Lemma 1 in their paper). We infer that any formula with less than $x_0 \cdot n/K$ constraints is satisfiable, therefore the same is true for any formula with $x_0 \cdot n/K$ clauses picked up from the clausal representation of constraints in Φ.

The second condition (expansion of the formula hypergraph) can be proved similarly.

□

6.5 Proof of Theorem 3

Proof.

1. For any real $r \geq 1$, formula F and set of clauses $G \subseteq F$, define the r-*deficiency of* G, $\delta_r(G) = r|Clauses(G)| - |Vars(G)|$. Also define

$$\delta_r^*(F) = \max\{\delta_r(G) : \emptyset \neq G \subseteq F\} \tag{5}$$

We claim that for any minimally unsatisfiable F, $\delta_{2k-3}^*(F) \geq 0$. Indeed, assume this was not true. Then there exists such F such that:

$$\delta_{2k-3}(G) \leq -1 \text{ for all } \emptyset \neq G \subseteq F. \tag{6}$$

Proposition 3 *Let F be a formula for which condition 6 holds. Then there exists an ordering $C_1, \ldots, C_{|F|}$ of constraints in F such that each constraint C_i contains at least $k-2$ variables that appear in* no *C_j, $j < i$.*

Proof. Denote by v_i the number of variables that appear in *exactly* i constraints of F. We have $\sum_{i\geq 1} i \cdot v_i = k \cdot |Constraints(F)|$, therefore $2|Var(F)| - v_1 \leq k \cdot |Constraints(F)|$. This can be rewritten as $v_1 \geq 2|Var(F)| - k|Constraints(F)| > |Constraints(F)| \cdot (2k-3-k) = (k-3) \cdot |Constraints(F)|$ (we have used the upper bound on $c^*(F)$. Therefore there exists at least one constraint in F with at least $k-2$ variables that are free in F. We set $C_{|F|} = C$ and apply this argument recursively to $F \setminus C$. $\square$

Call the $k-2$ new variables of C_i *free in* C_i. Call the other two variables *bound in* C_i. Let us show now that F cannot be minimally unsatisfiable. Construct a satisfying assignment for F incrementally: Consider constraint C_j. At most two of the variables in C_j are bound for C_j. Since C has no implicates of size at most two, one can set the remaining variables in a way that satisfies C_j. This yields a satisfying assignment for F, a contradiction with our assumption that F was minimally unsatisfiable.

Therefore $\delta_{2k-3}^*(F) \geq 0$, a statement equivalent to our conclusion.

2. To prove the resolution complexity lower bound we use the size-width connection for resolution complexity obtained in [7]: we prove that there exists $\eta > 0$ such that w.h.p. random instances of $SAT(\mathcal{C})$ having constraint density c have resolution width at least $\eta \cdot n$. We use the same strategy as in [7]

   (a) (prove that) w.h.p. minimally unsatisfiable subformulas are "large", and

   (b) any clause implied by a satisfiable formula of "intermediate" size contains w.h.p. "many" literals.

Indeed, define for a unsatisfiable formula Φ and (possibly empty) clause C

$$\mu(C) = \min\{|\Xi| : \Xi \subseteq \Phi, \Xi \models C\}.$$

Claim 2 *There exists $\eta_1 > 0$ such that for any $c > 0$, w.h.p. $\mu(\Box) \geq \eta_1 \cdot n$ (where Φ is a random instance of $SAT(\mathcal{C})$ having constraint density c).*

Proof. In the proof of Theorem 2 we have shown that there exists $\eta_0 > 0$ such that w.h.p. any unsatisfiable subformula of a given formula has at least $\eta_0 \cdot n$ constraints. Therefore *any* formula made of *clauses* in the CNF-representation of constraints in Φ, and which has less than $\eta_0 \cdot n$ clauses is satisfiable, and the claim follows, by taking $\eta_1 = \eta_0$. $\Box$

The only (slightly) nontrivial step of the proof, which critically uses the fact that constraints in $\mathcal{C}$ do not have implicates of length at most two, is to prove that clause implicates of subformulas of "medium" size have "many" variables. Formally (and proved in the Appendix)

Claim 3 *There exists $d > 0$ and $\eta_2 > 0$ such that w.h.p., for every clause C such that $d/2 \cdot n < \mu(C) <= dn$, $|C| \geq \eta_2 \cdot n$.*

Proof. Take $0 < \epsilon$. It is easy to see that if $c^*(F) < \frac{2}{2k-3+\epsilon}$ then w.h.p. for every subformula G of F, at least $\frac{\epsilon}{3} \cdot |Constraints(G)|$ have at least $k-2$ private variables: Indeed, since $c^*(G) < \frac{2}{2k-3+\epsilon}$, by a reasoning similar to the one we made previously $v_1(G) \geq (k-3+\epsilon)|Constraints(G)|$. Since constraints in G have arity k, at least $\epsilon/3 \cdot |Constraints(G)|$ have at least $k-2$ "private" variables.

Choose $y = \frac{2}{2k-3+\epsilon}$ in Lemma 6.1 for $\epsilon > 0$ a small enough constant. Since the problem has a sharp threshold in the region where the number of clauses is linear, $d = \inf\{x(y, c) : c >= c_{SAT(\mathcal{C})}\} > 0$. W.h.p. all subformulas of Φ having size less than $d/k \cdot n$ have a formula hypergraph that is (x, y)-sparse, therefore fall under the scope of the previous argument. Let Ξ be a subformula of Φ, having minimal size, such that $\Xi \models C$. We claim:

Claim 4 *For every clause P of Ξ with $k-2$ "private" variables, (i.e., one that does not appear in any other clause), at least one of these "private" variables appears in C.*

Indeed, suppose there exists a clause D of Ξ such that none of its private variables appears in C. Because of the minimality of Ξ there exists an assignment F that satisfies $\Xi \setminus \{D\}$ but does not satisfy D or C. Since D has no implicates of size two, there exists an assignment G, that differs from F only on the private variables of D, that satisfies Ξ. But since C does not contain any of the private variables of D, F coincides with G on variables in C. The conclusion is that G does not satisfy C, which contradicts the fact that $\Xi \models C$.

The proof of Claim 3 (and of item 2. of Theorem 3) follows: since for any clause K of one of the original constraints $\mu(K) = 1$, since $\mu(\Box) > \eta_1 \cdot n$ and since w.l.o.g. $0 < d < \eta_1$ (otherwise replace d with the smaller value) there exists a clause C such that

$$\mu(C) \in [d/2k \cdot n, d/k \cdot n]. \tag{7}$$

Indeed, let C' be a clause in the resolution refutation of Φ minimal with the property that $\mu(C') > dn$. Then at least one clause C, of the two involved in deriving C' satisfies equation 7.

By the previous claim it C contains at least one "private" variable from each clause of Ξ. Therefore $|C| \geq \eta_2 \cdot n$, with $\eta_2 = d/2k \cdot \epsilon$.

$\Box$