Journal of Universal Computer Science, vol. 13, no. 11 (2007), 1655-1670
submitted: 10/9/07, accepted: 19/11/07, appeared: 28/11/07 © J.UCS

Satisfying Assignments of Random Boolean Constraint Satisfaction Problems: Clusters and Overlaps[1]

Gabriel Istrate
(eAustria Research Institute
Bd. Vasile Pârvan 4, cam. 045B, Timişoara, RO-300223, Romania
Email: `gabrielistrate@acm.org`)

Abstract: The distribution of overlaps of solutions of a random constraint satisfaction problem (CSP) is an indicator of the overall geometry of its solution space. For random k-SAT, nonrigorous methods from Statistical Physics support the validity of the one step replica symmetry breaking approach. Some of these predictions were rigorously confirmed in [Mézard et al. 2005a] [Mézard et al. 2005b]. There it is proved that the overlap distribution of random k-SAT, $k \geq 9$, has discontinuous support. Furthermore, Achlioptas and Ricci-Tersenghi [Achlioptas and Ricci-Tersenghi 2006] proved that, for random k-SAT, $k \geq 8$, and constraint densities close enough to the phase transition:

- there exists an exponential number of clusters of satisfying assignments.
- the distance between satisfying assignments in different clusters is linear.

We aim to understand the structural properties of random CSP that lead to solution clustering. To this end, we prove two results on the cluster structure of solutions for binary CSP under the random model from [Molloy 2002]:

1. For all constraint sets S (described in [Creignou and Daudé 2004, Istrate 2005]) such that $SAT(S)$ has a sharp threshold and all $q \in (0, 1]$, q-overlap-$SAT(S)$ has a sharp threshold. In other words the first step of the approach in [Mézard et al. 2005a] works in all nontrivial cases.
2. For any constraint density value $c < 1$, the set of solutions of a random instance of 2-SAT form with high probability a single cluster. Also, for and any $q \in (0, 1]$ such an instance has with high probability two satisfying assignment of overlap $\sim q$. Thus, as expected from Statistical Physics predictions, the second step of the approach in [Mézard et al. 2005a] fails for 2-SAT.

Key Words: sharp thresholds, random constraint satisfaction, overlaps.
Category: G.2.1, G.3.

1 Introduction

A great deal of insight in the complexity of random constraint satisfaction problems has come from studying *phase transitions* [Monasson and Zecchina 1997]. Concepts from Statistical Physics, such as *first-order phase transitions*, *backbones*, or *replica symmetry breaking* have helped to refine (and understand the limitations of) the empirical observation that the "hardest" instances are located

[1] C. S. Calude, G. Stefanescu, and M. Zimand (eds.). *Combinatorics and Related Areas. A Collection of Papers in Honour of the 65th Birthday of Ioan Tomescu.*

at the transition point. In some cases the connection predicted by Statistical Physics can be made explicit in purely combinatorial terms. For instance Monasson et al. [Monasson et al. 1999b, Monasson et al. 1999a] have suggested that *first-order phase transitions* are correlated with exponential complexity of Davis-Putnam algorithms on random *unsatisfiable* instances at the phase transition. This has been rigorously confirmed to a certain extent [Achlioptas et al. 2001b, Achlioptas et al. 2004, Istrate et al. 2005]. For instances in the satisfiable phase, much of the intuition on the complexity of such instances comes again from Statistical Physics, via the so-called *one step replica symmetry breaking* (1-RSB) approach. The 1-RSB approach provides predictions on the geometric structure of the set of satisfying assignments of a random formula on n variables. The set of such assignments can be naturally viewed as a subgraph of the *hypercube of dimension* n, where two satisfying assignments are *neighbors* if they only differ in the value of one variable. Physics considerations imply that for small values of the constraint density c the set of satisfying assignments forms a single cluster. The distribution of overlaps is peaked around a certain constant value. The range of possible overlaps (even those that are exponentially infrequent) is a continuous interval. In the presence of 1-RSB, for constraint density values higher than a critical value c_{RSB} (smaller than the unsatisfiability threshold c_{UNSAT}) the set of satisfying assignments splits into several *clusters* such that: (i) assignments in the same cluster all agree on a set of variables having linear size. The distribution of overlaps of assignments in the same cluster is still concentrated around a constant; (ii) assignments in different clusters differ in $\Omega(n)$ variables; (iii) the distribution of all overlaps has discontinuous support (see Fig. 1 (a); note that recent studies [Krzakala et al. 2007] suggest the existence of further phases below c_{RSB}, omitted for simplicity from discussion and the figure). The geometry of satisfying assignments outlined above has implications for the complexity of heuristics such as local search, algorithms such as belief propagation, or Davis-Putnam. The 1-RSB approach provides (nonrigorous) values for the location of the phase transition in random k-SAT [Mertens et al. 2006] that seems to match the experimental evidence. Algorithms that take advantage of the geometry of solution space predicted by 1-RSB (e.g. the celebrated *survey propagation* algorithm [Braunstein et al. 2005]) have greatly extended the range of instances that can be solved in practice.

Rigorous results on the cluster structure of solutions of random CSP are emerging: Mézard et al. [Mézard et al. 2005a] have developed an ingenious method for proving that the distribution of overlap values of random k-SAT, with $k \geq 9$ indeed has discontinuous support. Their approach is based on the following concepts:

Definition 1. The *overlap* of two assignments A and B for a formula Φ on n variables, denoted by $overlap(A, B)$, is the fraction of variables on which the

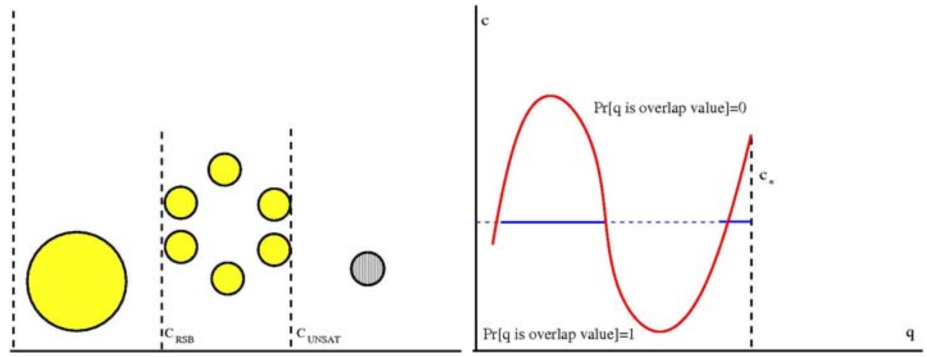

Figure 1: (a) Structure of solution space according to 1RSB predictions.(b) Graphical description of the method used in [Mézard et al. 2005a] to prove the discontinuity of support of the overlap distribution.

two assignments agree (this is similar to [Mézard et al. 2005a] and linearly related to the notion of overlap from the statistical physics literature, where truth values are modeled by $+1$ and -1, instead of 0/1). Formally $overlap(A, B) = \frac{|\{i: A(x_i)=B(x_i)\}|}{n}$.

The distribution of overlaps is, indeed, the original order parameter that was originally used to study the phase transition in random k-SAT, see the paper [Monasson and Zecchina 1997].

Definition 2. **q-overlap-k-SAT**: Given a k-CNF formula Φ on n variables, decide whether Φ has two satisfying assignments A and B such that $overlap(A, B) \in [q - 1/\sqrt{n}, q + 1/\sqrt{n}]$ (following the suggestion in [Mézard et al. 2005a], we will use the function $1/\sqrt{n}$ for the the width of the possible overlap around q; as discussed there, similar results are obtained with any "reasonably large" function $b(n) = o(n)$). We will refer to this event as *A and B have overlap approximately equal to q*.

For every value of q, the probability that a random k-SAT formula has two assignments with overlap $\sim q$ is monotonically decreasing with constraint density, and is empirically changing from 1 to 0 around a critical value $c_{k,q}$ of the constraint density. If one can show that the function $W : q \to c_{k,q}$ is *not* monotonic then *there exists a critical value c_* such that the horizontal line at c_* will intersect the graph of function W at multiple points. Therefore (Figure 1 (b)) the distribution of overlaps in a random k-SAT formula of constraint density c_* has discontinuous support* (these results were further extended, for k-SAT, $k \geq 9$, by Achlioptas and Ricci-Tersenghi [Achlioptas and Ricci-Tersenghi 2006]).

Our ultimate goal is to obtain an understanding of the underlying reasons for the emergence of clustering in random CSP, with an attempt at a precise classification. We investigate the nature of overlap distributions of CSP under the random model defined and investigated by Molloy [Molloy 2002]. We cannot obtain a complete classification (whether the results in [Mézard et al. 2005a, Achlioptas and Ricci-Tersenghi 2006] extend to random 3-SAT is a more subtle problem; see [Maneva and Sinclair 2007]). Instead, we prove two partial results: Theorem 8 shows that the first step of Mézard's approach can be applied to all random CSP problems with a sharp threshold. In contrast, in Theorem 9 we show that satisfying assignments of random instances of 2-SAT in the satisfiable phase form a single cluster, and can yield all possible values of the overlap. This confirms the prediction [Monasson and Zecchina 1997] that the solutions space of 2-SAT has a different nature, describes by the so-called "replica symmetric" approach. The two results above are also naturally related to results of Gopalan et al. [Gopalan et al. 2006]. They proved a dichotomy theorem for the complexity of deciding whether the set of satisfying assignments of a CSP is connected (under the usual notion of adjacent assignments). One ingredient of the result is a restriction (called *tightness*) on the nature of constraints involved. Theorem 2 provides a natural examples of CSP with tight constraints for which there is evidence (the continuity of the overlap distribution) that symmetry breaking does not take place. It also shows that, to be really meaningful, the definition of adjacent assignments from [Gopalan et al. 2006] should be somewhat modified.

2 Preliminaries

Throughout the paper we will assume familiarity with the general concepts of phase transitions in combinatorial problems (see e.g. [Martin et al. 2001]) and random structures. One paper whose concepts and methods we use in detail (and we assume greater familiarity with) is [Friedgut 1999].

Consider a monotonically increasing problem $A = (A_n)$ under the *constant probability model* $\Gamma(n, p)$. For $\epsilon > 0$ let $p_\epsilon = p_\epsilon(n)$ define the canonical probability such that $\Pr_{x\in\Gamma(n,p_\epsilon(n))}[x \in A] = \epsilon$. The probability that a random sample x satisfies property A (i.e. $x \in A$) is a monotonically increasing function of p.

Definition 3. Problem A has a *sharp threshold* iff for every $0 < \epsilon < 1/2$, we have

$$lim_{n\to\infty} \frac{p_{1-\epsilon}(n) - p_\epsilon(n)}{p_{1/2}(n)} = 0.$$

A has *a coarse threshold* if for some $\epsilon > 0$ it holds that

$$\underline{\lim}_{n\to\infty} \frac{p_{1-\epsilon}(n) - p_\epsilon(n)}{p_{1/2}(n)} > 0.$$

Related definitions can be given for the other two models for generating random structures, the *counting model* and *the multiset model* [Bollobás 1985]. Under reasonable conditions [Bollobás 1985] these models are equivalent, and we will liberally switch between them. In particular, for satisfiability problem A, and an instance Φ of A, $c_A(\Phi)$ will denote its *constraint density*, the ratio between the number of clauses and the number of variables of Φ. To specify the random model in this latter cases we have to specify the constraint density as a function of n, the number of variables. We will use c_A to denote the value of the constraint density $c_A(\Phi)$ (in the counting/multiset models) corresponding to taking $p = p_{1/2}$ in the constant probability model. c_A is a function on n that is believed to tend to a constant as $n \to \infty$. However, Friedgut's proof [Friedgut 1999] of a sharp threshold in k-SAT (and our results) leave this issue open.

Definition 4. Let $\mathcal{D} = \{0, 1, \ldots, t-1\}$, $t \geq 2$ be a fixed set. Consider the set of all $2^{t^k} - 1$ potential nonempty binary constraints on k variables $X_1, \ldots, X_k$. We fix a set of constraints $\mathcal{C}$ and define the random model $CSP(\mathcal{C})$. A random formula from $CSP_{n,p}(\mathcal{C})$ is specified by the following procedure:

(i) n is the number of variables.

(ii) for each k-tuple of *ordered* distinct variables $(x_1, \ldots, x_k)$ and each $C \in \mathcal{C}$ add constraint $C(x_1, \ldots, x_k)$ independently with probability p.

We will write $SAT(\mathcal{C})$ instead of $CSP(\mathcal{C})$ for *boolean* constraint satisfaction problems (i.e. $t = 2$).

Definition 5. Let $\mathcal{D} = \{0, 1, \ldots, t-1\}$, $t \geq 2$ be a fixed set. Let q be a real number in the range [0,1]. The problem q-overlap-$CSP(\mathcal{C})$ is the decision problem specified as follows:

(i) The input is an instance Φ of $CSP_{n,p}(\mathcal{C})$.

(ii) The decision problem is whether Φ has two satisfying assignments A, B such that $overlap(A, B) \in [q - 1/\sqrt{n}, q + 1/\sqrt{n}]$ (following [Mézard et al. 2005b], we will informally refer to the property as *"Φ is q-satisfiable"*).

The random model for q-overlap-$CSP(\mathcal{C})$ is simply the one for $CSP_{n,p}(\mathcal{C})$. We will refer to this class of problems as *fixed-overlap CSP*.

The notion of adjacent satisfying assignments used by Achlioptas and Ricci-Tersenghi in [Achlioptas and Ricci-Tersenghi 2006], while adequate for random k-SAT, is not suited for other random CSP. For instance, it is impossible to flip exactly one bit in a satisfying assignment of an instance of 1-in-k SAT [Achlioptas et al. 2001a] and still obtain a satisfying assignment (except for the case when that variable does not appear in the formula). Thus we will use the following setup: let $f(n) = o(n)$ be a suitably large function; we will assume that $\lim f(n)/\log n = \infty$. Two satisfying assignments that differ on at most

$f(n)$ variables will be called adjacent. A *cluster* is a connected component of the set of satisfying assignments.

3 Results

In this section we study the sharpness of the threshold for random generalized constraint satisfaction problem defined by Molloy [Molloy 2002]. Creignou and Daudé [Creignou and Daudé 2004] (and independently the author of this paper [Istrate 2005]) have characterized the boolean CSP problems $SAT(\mathcal{C})$ with a sharp threshold:

Definition 6. A set of constraints $\mathcal{C}$ is *interesting* if there exist constraints $C_0, C_1 \in \mathcal{C}$ with $C_0(\overline{0}) = C_1(\overline{1}) = 0$, where $\overline{0}, \overline{1}$ are the "all zeros" ("all ones") assignments. Constraint C_2 is *an implicate of* C_1 iff every satisfying assignment for C_1 satisfies C_2. A boolean constraint C *strongly depends on a literal* if it has an unit clause as an implicate. A boolean constraint C *strongly depends on a 2-XOR relation* if $\exists i, j \in \{1, \ldots, k\}$ such that constraint "$x_i \neq x_j$" is an implicate of C.

Proposition 7. *[Creignou and Daudé 2004, Istrate 2005] Consider a generalized satisfiability problem $SAT(\mathcal{C})$ with $\mathcal{C}$ interesting. (i) If some constraint in $\mathcal{C}$ strongly depends on one literal then $SAT(\mathcal{C})$ has a coarse threshold; (ii) If some constraint in $\mathcal{C}$ strongly depends on a 2XOR-relation then $SAT(\mathcal{C})$ has a coarse threshold; (iii) In all other cases $SAT(\mathcal{C})$ has a sharp threshold.*

Mora et. al [Mézard et al. 2005b] proved that all problems q-overlap-k-SAT, $k \geq 2$ have a sharp threshold. We extend this result by showing that for all CSP with a sharp threshold, their fixed-overlap versions also have a sharp threshold:

Theorem 8. *Consider a generalized satisfiability problem $SAT(\mathcal{C})$ such that (i) $\mathcal{C}$ is interesting (ii) No constraint in $\mathcal{C}$ strongly depends on a literal; (iii) No constraint in $\mathcal{C}$ strongly depends on a 2XOR-relation. Then for all values $q \in (0, 1]$ the problem q-overlap-$SAT(\mathcal{C})$ has a sharp threshold.*

The previous result does not yet rigorously prove the existence of curve W since it does not prove fact that the phase transition in the q-overlap versions happens at some constant constraint density c_q .

Given the previous result, how can a problem $SAT(\mathcal{C})$ have an overlap distribution with continuous support ? Obviously, the second step of the approach in [Mézard et al. 2005a] must fail. This happens when *the location c_q of the transition for the q-overlap version of $SAT(\mathcal{C})$ is a monotonic function of the overlap q*. The next result shows gives a natural problem for which this is indeed the case:

Theorem 9. *The following are true:*

(i) Let $c < 1$. Then with probability $1 - o(1)$ the satisfying assignments of a random instance of 2-SAT of constraint density c form a single cluster.

(ii) Also, let $q \in (0, 1]$. Let $c < 1$. Then with probability $1 - o(1)$ a random instance of 2-SAT of constraint density c is q-satisfiable.

4 Proof of Theorem 8

Before presenting the proof, let us remark that for boolean constraints, the hypothesis of the Theorem 8 implies that the set of constraints $\mathcal{C}$ is *well-behaved*. That is[Molloy 2002], every formula whose hypergraph is tree-like or unicyclic is satisfiable. This is, for instance, an easy consequence of conditions (D0),(D1), Theorem 4.1 in [Creignou and Daudé 2004]. Also, since $\mathcal{C}$ is interesting there exist constraints $\Gamma_0, \Gamma_1 \in \mathcal{C}$ such that $\Gamma_0(x_1, \ldots, x_k) \models \overline{x_1} \vee \ldots \vee \overline{x_k}$ and $\Gamma_1(x_1, \ldots, x_k) \models x_1 \vee \ldots \vee x_k$.

We will employ the Friedgut-Bourgain criterion for the existence of a sharp threshold of a monotonic property A. Note that any problem q-overlap-SAT($\mathcal{C}$) is indeed monotone, since adding clauses can only reduce the set of satisfying assignments, in particular decreasing the probability of q-satisfiability. The starting point of all applications of the Friedgut-Bourgain criterion is noting that if a monotone property A has a coarse threshold then there exists $0 < \epsilon < 1/2$, $p^* = p^*(n) \in [p_{1-\epsilon}, p_\epsilon]$ and $C > 0$ such that $p \cdot \frac{d\mu_p(A)}{dp}|_{p=p^*(n)} < C$. Bourgain and Friedgut have shown that the following holds:

Proposition 10. *Suppose $p = o(1)$ is such that $p \cdot \frac{d\mu_p(A)}{dp}|_{p=p^*(n)} < C$. Then there is $\delta = \delta(C) > 0$ such that either $\mu_p(x \in \{0,1\}^n | \ x$ contains $x' \in A$ of size $|x'| \leq 10C\} > \delta$, or there exists $x' \notin A$ of size $|x'| \leq 10C$ such that $\mu_p(x \in A | x \supset x') > \mu_p(A) + \delta$.*

(In fact, in [Friedgut 1999] the proposition is stated assuming for convenience that $p = p_{1/2}$, but this is not needed. We give here the general statement). We will need, in fact, an enhancement to the Bourgain-Friedgut result that was given by Friedgut in [Friedgut 2005]: For a finite set of words W define *the filter generated by* W, $F(W)$ as $F(W) = \{x \mid (\exists y \in W) \text{ with } x \supseteq y\}$. Friedgut noted ([Friedgut 2005], remarks on pages 5-6 of that paper) that the set W of "booster" sets x' in the second conditions satisfies $\mu_p(F(W)) = \Omega(1)$.

Consider now a set of constraints $\mathcal{C}$ satisfying the conditions the Theorem, and let $A = \overline{\text{q-overlap-SAT}(\mathcal{C})}$. Applying Proposition 10 enhanced by the previous observation, and taking into account the fact that the number of isomorphism types of formulas of size at most $10C$ is finite, we infer that we can assume that formula x' in the second condition appears with probability $\Omega(1)$ as a subformula in a random formula in q-overlap-$\mathrm{SAT}_p(\mathcal{C})$. Furthermore, instead on conditioning on the presence of x' as a subset of x one can, instead, add it. Finally, note that for random constraint satisfaction problems, because of the invariance of such problems under variable renaming, one only needs to add a random copy of x'. Putting all these observations together, the following version of Proposition 10 holds:

Proposition 11. *Suppose $p = o(1)$ is such that $p \cdot \frac{d\mu_p(A)}{dp}|_{p=p^*(n)} < C$. Then there is $\delta = \delta(C) > 0$ such that either*

$$\mu_p(x \in \{0,1\}^n | \ x \textit{ contains } x' \in A \textit{ of size } |x'| \leq 10C\} > \delta \tag{1}$$

or there exists $F \notin A$ of size $|F| \leq 10C$, such that

- *Formula F appears with probability $\Omega(1)$ as a subformula in a random formula in $CSP_p(\mathcal{C})$.*
- *If Ξ denotes the formula obtained by creating a copy of x' on a random set of variables, then*

$$\mu_p(x \cup \Xi \in A) > \mu_p(A) + \delta. \tag{2}$$

To show that random q-overlap-SAT$(\mathcal{C})$ has a sharp threshold, we will reason by contradiction. Assuming this is not the case, one needs to prove that the two conditions in Proposition 11 do not hold.

Suppose, indeed, that condition (1) was true. That is, with positive probability it is true that a random formula $\Phi \in CSP(\mathcal{C})$ contains some subformula $\Phi' \in \overline{\text{q-overlap-SAT}(\mathcal{C})}$ of size at most $10C$. With high probability all subformulas of a random formula Φ of size at most $10C$ are either tree-like or unicyclic. But because the set of constraints $\mathcal{C}$ is well-behaved (this is the point where the hypothesis on the constraint set $\mathcal{C}$ is used), all formulas in $CSP(\mathcal{C})$ that are tree-like or unicyclic are satisfiable. Since the formula contains a finite number of variables, one can set the other variables not appearing in Φ in a way that will create two satisfying assignments with overlap approximately q. Therefore the first condition in Proposition 11 cannot be true.

Assume, now, that condition (2) is true. The condition means there exists $F \in q$-overlap-SAT$(\mathcal{C})$, a formula of size at most $10C$, such that adding F to a random formula $\Phi \in CSP_p(\mathcal{C})$ diminishes the probability that the resulting formula has two assignments of overlap $\simeq q$ by at least a constant δ. As discussed, we assume that F occurs with probability $\Omega(1)$ in a random formula in $CSP_p(\mathcal{C})$. Therefore F is tree-like or unicyclic.

Definition 12. A *unit clause* is a constraint (not necessarily part of the constraint set $\mathcal{C}$) specified by a condition $X = \delta$, with X being a variable and $\delta \in \{0,1\}$.

Lemma 13. *If F satisfies condition (2) then there exists another formula G that is specified by a finite conjunction of unit clauses $G \equiv (X_1 = \delta_1) \wedge \ldots \wedge (X_p = \delta_p)$, that also satisfies condition (2).*

Proof. Formula F appears with constant probability in a random $CSP(\mathcal{C})$ formula with probability p and has constant size. Therefore F is either tree-like or unicyclic. The result follows easily by replacing F with formula G consisting of the conjunction of unit constraints corresponding to a satisfying assignment of F. Indeed, G is tighter than F, so adding a random copy of G instead of a random copy of F can only increase the probability that the resulting formula is unsatisfiable. □

The key to refuting condition (2) is to show that, if it did hold then, for every monotonically increasing function $f(n)$ that tends to infinity, we could also increase the probability of unsatisfiability by a positive constant if, instead of conditioning on x containing a copy of F, we add $f(n)$ random constraints from set $\mathcal{C}$. We first prove:

Lemma 14. *Let $0 < \tau < 1$ be a constant and let p be such that $\mu_p(q-overlap-SAT(\mathcal{C})) \geq \tau$. Assume that $r \geq 1$ and that $g_1, g_2, \ldots g_r$ are elements of $\{0,1\}$ such that, when $(X_1, X_2, \ldots, X_r)$ is a random r-tuple of different variables*

$$Pr(\Phi \text{ has sat. assign. } A, B \text{ of overlap } \simeq q \text{ with } X_1 = g_1, \ldots, X_r = g_r) \leq \frac{\tau}{2}. \tag{3}$$

Then there exists constant $m \geq 1$ (that only depends on k, r, τ) such that, if η denotes a formula from $CSP(\mathcal{C})$ obtained by adding, for each $x \in \{0,1\}$, $m \cdot r \cdot 2^{k^r}$ random copies of Γ_x, then

$$Pr(\Phi \cup \eta \in q\text{-overlap-}SAT(\mathcal{C})) \leq \frac{\tau}{2} \tag{4}$$

Proof.

For $i \in \{1, \ldots, r\}$ define A_i to be the event that the formula Φ has a pair of satisfying assignments of overlap $\simeq q$ with $X_1 = g_1, \ldots, X_i = g_i$. Also define A_0 to be the event that $\Phi \in$ q-overlap-SAT($\mathcal{C}$). The hypothesis translates as the fact that both inequalities $Pr(A_0) \geq \tau$ and $Pr(A_r) \leq \frac{\tau}{2}$ are true. Therefore $Pr(A_r|A_0) = \frac{Pr(A_r \wedge A_0)}{Pr(A_0)} \leq \frac{\tau/2}{\tau} = \frac{1}{2}$. Since $\overline{A_{r-1}} \Rightarrow \overline{A_r}$ we have

$$\mu_r := Pr[\overline{A_r}|A_0] = Pr[\overline{A_{r-1}}|A_0] + Pr[\overline{A_r}|A_{r-1} \wedge A_0] \cdot Pr[A_{r-1}|A_0] \geq \frac{1}{2} \tag{5}$$

But $Pr[\overline{A_r}|A_{r-1} \wedge A_0] = Pr[\overline{A_r}|A_{r-1}]$ is the fraction of variables in formula $\Phi \wedge (X_1 = g_1) \wedge \ldots \wedge (X_{r-1} = g_{r-1})$ that have to receive values different from g_r in order for the resulting formula to still have two satisfying assignments of overlap $\sim q$; let C_r be the set of such variables. If instead of the last unit constraint we add a random copy of constraint Γ_{g_r}, the resulting formula is in

$\overline{\text{q-overlap-SAT}(\mathcal{C})}$ when all the variables appearing in the new constraint are in the set C_r. Denoting $\lambda_r = Pr[\overline{A_r}|A_{r-1}]$, the probability of this last event happening is $\lambda_r^k/(1-o(1))$ (we choose a k-tuple of distinct variables from a set of density λ_r); Thus the probability that the new formula is in $\overline{\text{q-overlap-SAT}(\mathcal{C})}$ is at least $\nu_r := Pr[A_{r-1}|\overline{A_0}] + \frac{\lambda_r^k}{1-o(1)} \cdot Pr[\overline{A_{r-1}}|\overline{A_0}]$. Applying Jensen's inequality to the convex function $f(x) = x^k$ and using inequality (5), we infer

$$\begin{aligned} \frac{1}{2^k} \le \mu_r^k &= (Pr[A_{r-1}|\overline{A_0}] \cdot 1 + Pr[A_r|\overline{A_{r-1}}] \cdot Pr[\overline{A_{r-1}}|\overline{A_0}])^k \le \\ &\le Pr[A_{r-1}|\overline{A_0}] \cdot 1^k + Pr[A_r|\overline{A_{r-1}}]^k \cdot Pr[\overline{A_{r-1}}|\overline{A_0}] = \\ &= (Pr[A_{r-1}|\overline{A_0}] + \lambda_r^k \cdot Pr[\overline{A_{r-1}}|\overline{A_0}]) = \nu_r \cdot (1+o(1)). \end{aligned}$$

Thus $\nu_r \ge \frac{1}{2^k} \cdot (1-o(1))$. The conclusion of this long argument is that adding one random copy of Γ_{b_r} instead of the r-th constraint lowers the probability of membership to q-overlap-SAT$(\mathcal{C})$ to no less than $\frac{1}{2^k} \cdot (1-o(1))$. Adding the copy of the constraint *before* the first $r-1$ unit constraints and repeating the argument recursively implies the fact that, if instead of adding the r unit constraints to Φ we add r random copies of $\Gamma_{b_1}, \ldots, \Gamma_{b_r}$ that the resulting formula belongs to $\overline{\text{q-overlap-SAT}(\mathcal{C})}$, given that $\Phi \in q$-overlap-SAT$(\mathcal{C})$, is at least $\gamma_r = \frac{1}{2^{k^r}(1-o(1))}$. Since the values $b_1, \ldots, b_r$ can repeat themselves, the same is true if we add r random copies of Γ_x for *every* x.

Suppose now that we add $r \cdot m \cdot 2^{k^r}$ copies of each Γ_x (that is, we repeat the random experiment $m \cdot 2^{k^r}$ times, for some integer $m \ge 1$). The probability that none of the experiments will make the resulting formula unsatisfiable is at most $(1-\gamma_r)^{m \cdot 2^{k^r}}$. For some constant m this is going to be at most $1 - \frac{\tau}{2}$. This means that $Pr(\Phi \cup \eta$ is satisfiable$) \le \frac{\tau}{2}$. □

We can refute condition (4) directly, thus obtaining a contradiction. To do so, we employ the following result (Lemma 3.1 in [Achlioptas and Friedgut 1999]):

Lemma 15. *For a monotone property*[2] A *let* $\mu(p) = Pr[G \in \Gamma(n,p)$ *has property* $A]$, *and let* $\mu^+(p, M) = Pr[G_1 \cup G_2 \mid G_1 \in \Gamma(n,p), G_2 \in \Gamma(n, M)$ *has property* $A]$. *Let* $A = A(n) \subseteq \{0,1\}^n$ *be a monotone property and* $M = M(n)$ *such that* $M = o(\sqrt{np})$. *Then* $|\mu(p) - \mu^+(p, M)| = o(1)$.

We obtain a contradiction in the following way: consider a random formula η with $f(n)$ clauses, for some $f(n) \to \infty$. It is easy to show that the probability that η contains, for some x, less than $r \cdot m \cdot 2^{k^r}$ copies of Γ_x (with r, m as in Lemma 2) is $o(1)$. So adding η (instead of the random formula in Lemma 14) decreases the probability of q-satisfiability by at least $\delta - o(1)$. But this contradicts the conclusion of Lemma 15. □

[2] Achlioptas and Friedgut assume A to be a monotone *graph* property, but this fact is not used anywhere in their proof.

5 Proof of Theorem 9

We will use the well-known graph-theoretic interpretation of 2-CNF formulas, that associates to a given formula Φ on n variables a directed graph G_Φ with $2n$ vertices $\{x_1, \ldots, x_n, \overline{x_1}, \ldots, \overline{x_n}\}$, and for every clause $C = \alpha \vee \beta$ of Φ it adds directed edges $\overline{\alpha} \to \beta$ and $\overline{\beta} \to \alpha$ to G_Φ. We will need a number of results from [Rozenthal et al. 1999] concerning the structure of graph G_Φ when Φ is a random formula of constraint density $c < 1$.

Definition 16. A *cycle* is a set $l_1 \to l_2, l_2 \to l_3, \ldots, l_s \to l_1$ of directed edges. Two cycles C_1, C_2 are *overlapping* if they share at least an edge. Two cycles C_1, C_2 are *connected by a path* if there exist vertices $x \in C_1, y \in C_2$ and a path (possibly of length zero, i.e. $x = y$) from x to y.

Lemma 17. *Let $t = t(n)$ such that $1 = o(t)$. Let Φ be a random 2-CNF formula of constraint density $c < 1$ and G_Φ be its associated digraph. With probability $1 - o(1)$ the following are true: (i) G_Φ contains no cycles connected by a path. (ii) G_Φ contains no overlapping cycles. (ii) the sum of all the cycle lengths is less than t.*

To these results we add the following claim (whose proof is similar to that of Claim 17 (i) from [Rozenthal et al. 1999]): With probability $1 - o(1)$ no literal implies literals in two different cycles.

We can thus divide the literals of the formula into four classes: (i) those that are on a cycle. (ii) those that are not on a cycle, but *imply* a literal on a cycle. (iii) those that are not on a cycle, but are *implied* by a literal on a cycle. (iv) those that are not on a cycle and neither imply nor are implied by a literal on a cycle.

Definition 18. A literal x is *bad* if there exists y such that $x \xrightarrow{*} y$, $x \xrightarrow{*} \overline{y}$.

We first claim that there is a function $h(n) = o(n)$ such that with probability $1 - o(1)$ the number of bad literals is at most $h(n)$. Indeed, all bad literals can only be set to false in any satisfying assignment of the formula. This means that a bad literal belongs to the *spine* of the formula [Bollobás et al. 2001]. But a standard argument (see e.g. [Istrate et al. 2005]) shows that the size of the spine is $o(n)$.

Bad literals (and their negations) are assigned fixed values in all satisfying assignments. This property guarantees that such literals do not influence the value of the overlap between any two satisfying assignments. Let $\mathcal{B}$ be the set of such literals.

Theorem 9(i): Let A and B be two satisfying assignments of a formula Φ, such that $d(A, B) > \log n$ (i.e. A and B are not adjacent). We will prove the following result:

Lemma 19. *There exists a satisfying assignment C such that $d(A, C) = O(\log n)$ and $d(C, B) < d(A, B)$. That is, C is adjacent to A and closer to B than A.*

An iterative application of the lemma proves the Theorem 9(i).

Proof: Let x be a variable such that $A(x) \neq B(x)$ and x is *implication minimal* with this property. In other words if $y \neq x$ and $y \xrightarrow{*} x$ then $A(y) = B(y)$.
Case 1: $A(x) = 0$ and $B(x) = 1$. Then $B(z) = 1$ for all z such that $x \xrightarrow{*} z$. Define the assignment C by $C(z) = 1$ if $x \xrightarrow{*} z$, $C(z) = A(z)$ otherwise. It is clear that $d(C, B) < d(A, B)$, since C coincides with B on all bits whose value changes. To show that C is a satisfying assignment, suppose C did not satisfy some clause $W = (\alpha \vee \beta)$. Then one of the following is true.

(1): both α and β are negations of literals implied by x. This leads to a contradiction, since it would imply that B does not satisfy clause $\alpha \vee \beta$ either.

(2): one of them (say α) is the negation of a literal implied by x. Since $x \xrightarrow{*} \overline{\alpha}$ and $\overline{\alpha} \rightarrow \beta$, it follows that $C(\beta) = 1$, so C satisfies clause W.

(3): none of them is the negation of a literal implied by x. Then $C(\alpha) = A(\alpha)$ and $C(\beta) = A(\beta)$, a contradiction, since A satisfies clause W.

Case 2: $A(x) = 1$ and $B(x) = 0$. Then $B(z) = 0$ for all z such that $z \xrightarrow{*} x$. Define the assignment C by $C(z) = 0$ if $z \xrightarrow{*} x$, $C(z) = A(z)$ otherwise. It is clear that $d(C, B) < d(A, B)$, since C coincides with B on all the bits that change value, one of which is x. To show that C is a satisfying assignment, suppose C did not satisfy some clause $\alpha \vee \beta$. Then one of the following cases must hold

(1): both α and β are literals that imply x. This leads to a contradiction, since this would mean that B with respect to satisfying clause $\alpha \vee \beta$.

(2): one of them (say α) implies x. Since $\overline{\beta} \xrightarrow{*} \alpha$, it follows that $\overline{\beta} \xrightarrow{*} x$, therefore β is assigned the value TRUE by C, a contradiction.

(3): none of α, β implies x. Then C and A coincide with respect to the values they give to α, β, a contradiction, since A satisfies clause W.

Theorem 9(ii): We directly construct two satisfying assignments A and B of overlap $qn \pm \sqrt{n}$. We will work with a directed weighted graph G_2 obtained from G_Φ by contracting every cycle to a node and assigning this node a weight equal to twice the size of the contracted cycle. G_2 is well-defined when cycles in G_Φ do not intersect, an event that happens (cf. Claim 17) with probability $1 - o(1)$. All literals on a cycle of G_Φ need, of course, to be given the same value in any satisfying assignment. Since we have contracted all cycles in G_Φ, G_2 is a directed acyclic graph. The set of nodes corresponding to bad literals is downward closed, because if $x \rightarrow y$ and y is bad then x is bad. Correspondingly, the set of nodes corresponding to negations of a bad literal is upward closed.

We begin by defining a set S of nodes of G_2 that will ultimately contain half of the nodes in G_2. Nodes not chosen in S will be referred to as *eliminated*). In parallel we build a partial assignment by assigning those literals corresponding to eliminated nodes the unique values that are consistent with the satisfiability of the formula. Set S is recursively specified as follows:

(i) start by defining V to be the set of all nodes in G_2.

(ii) add all nodes of of indegree 0 in V to S and eliminate all nodes of outdegree 0. Set V to be the set of remaining nodes (not added to S or eliminated).

(iii) continue this process as long as $V \neq \emptyset$.

It is easy to see that the set of literals corresponding to nodes in S contains, for every variable x, exactly one of x and $\overline{x}$. Indeed, one cannot add both x and $\overline{x}$ to S in one step, otherwise the pure literal implying both would be bad. But then, when adding one of them we immediately eliminate the other one. On the other hand, we only eliminate a literal when its opposite has been retained in S.

The first assignment, A simply corresponds to setting all literals corresponding to nodes in S to TRUE. We define the second assignment iteratively by the following process:

(i) in Stage 1 choose a node of indegree zero, assign its associated variable the value FALSE and eliminate the node from S. If the eliminated node corresponds to a cycle in G_2 all variables in the cycle are set to FALSE.

(ii) when a remaining node becomes of indegree zero as a result of eliminations, it is labeled by the value of the stage that led to this happening (nodes that originally had indegree zero are labelled 0).

(iii) the literal chosen to set to FALSE is among those with a smallest stage number.

(iv) continue the process until the number of variables assigned FALSE is in the interval $[qn - \sqrt{n}, qn + \sqrt{n}]$. This is possibly if the sum of all cycle lengths in the formula graph of Φ is $o(\sqrt{n})$, which happens (cf. Lemma 17) with probability $1 - o(1)$.

(v) The remaining literals in S are set to TRUE.

Because bad literals are assigned identical values in both A and B it is easy to see that $overlap(A, B) \in [qn - \sqrt{n}, qn + \sqrt{n}]$. We complete the proof of Theorem 9 by:

Lemma 20. *A and B are satisfying assignments for Φ.*

Proof: Suppose there exists a clause $C = (\overline{x} \vee y) \equiv (x \rightarrow y)$ of Φ that is not satisfied by A. Then x is given a TRUE value and y is given a FALSE value.

Thus either $\overline{x}$ is a bad literal, or x is in S. Also, either y is a bad literal or $\overline{y}$ is in S. Suppose y were a bad literal. Then, since $x \rightarrow y$, x is also bad. But this contradicts the two possible alternatives ($\overline{x}$ is a bad literal or x is in S). Suppose now $\overline{y}$ is in S. Then $C \equiv (\overline{y} \rightarrow \overline{x})$. Therefore, either $\overline{x} \in S$ or $\overline{x}$ is among the literals (bad literals and their negations) eliminated before defining S. The first alternative leads to a contradiction with the two possible alternatives ($\overline{x}$ is a bad literal or x is in S), so it must be that $\overline{x}$ is a bad literal. But then $\overline{y}$ is also bad, contradicting the assumption that $\overline{y}$ is in S.

A similar argument shows that B is a satisfying assignment. Indeed, suppose there existed a clause $C = (\overline{x} \vee y) \equiv (x \rightarrow y)$ of Φ not satisfied by B. Then $B(x) = TRUE$, $B(y) = FALSE$. The choices compatible with this setup are: (i) x is in S and $B(x) = TRUE$, or $\overline{x}$ is bad. (ii) y is bad, or y is in S and $B(y) = FALSE$, or $\overline{y} \in S$ and $B(\overline{y}) = TRUE$, i.e. $B(y) = FALSE$. First, if y were bad then so would be x, contradicting all possible choices in (i). If x, y were both in S, with y assigned FALSE, since by construction of B the set of literals in S is downward closed under implication it follows that x would also be assigned FALSE, a contradiction. The other other possibility is that $\overline{x}$ is bad. But since $\overline{y} \rightarrow \overline{x}$ that would mean that $\overline{y}$ is bad, a contradiction with the assumption that $y \in S$. Finally, assume $\overline{y}$ is in S and is assigned TRUE. Since $\overline{y} \rightarrow \overline{x}$ either $\overline{x} \in S$ or x is a bad literal. In the first case, since the set of literals assigned to TRUE is upward closed under implication it would mean that $\overline{x}$ is assigned TRUE by B, i.e. x is assigned FALSE, a contradiction. Suppose now that x is bad. Then $B(x) = 0$, a contradiction. □

6 Conclusions

To sum up, we have analyzed the method introduced by Mora et al. for proving the existence of multiple clusters of solutions, and identified the factor that makes their approach fail: the second condition in their approach. We have also given an example of a problem where this happens, random 2-satisfiability.

Our results raise the question of characterizing all random constraint satisfaction problems that display solution clustering and discontinuity in their overlap distribution. Note that such general results exist for a couple of problems, such as the existence of a sharp threshold [Creignou and Daudé 2004], [Istrate 2005] or the discontinuity of the spine order parameter [Istrate et al. 2005]. In particular, in this latter problem random 2-satisfiability plays an important role, in that problems with a continuous spine seem to be "2-SAT-like".

We believe that this might be the case for the discontinuity of the overlap distribution as well. In particular, investigating the geometry of the solution space of random 1-in-k satisfiability [Achlioptas et al. 2001a, Raymond 2007, Maneva et al. 2007] is a natural continuation of this work. We have recently obtained some results (amounting to a "physicist's proof", but not yet a complete rigorous argument) that random 1-in-k SAT has a single cluster of solutions. The gist of this argument is a mapping of the set of solutions of a random 1-in-k CNF formula of subcritical density lying between two given solutions A and B to the set of maps from the vertices of a random graph $G(n, c/n)$, $c < 1$ to $\{0, 1\}$ that are constant on each connected component. A and B themselves correspond to the constant maps **0** and **1**. The fact that w.h.p. the largest connected component of the graph has size $O(\log n)$ means that A and B are connected by a sequence of satisfying assignments for the original formula such that two consecutive assignments differ in $O(\log n)$ positions (corresponding to changing the label on one connected component).

Whether the argument outlined above can be made completely rigorous and, more generally, whether all problems $SAT(\mathcal{C})$ with a single cluster of solutions is "2-SAT like" in some well-defined rigorous sense is an interesting research problem.

Acknowledgments

This work has been supported by a Marie Curie International Reintegration Grant within the 6th European Community Framework Programme.

References

[Achlioptas and Friedgut 1999] Achlioptas, D., Friedgut, E.: "A sharp threshold for k-colorability"; Random Structures and Algorithms 14, 1 (1999), 63-70.

[Achlioptas and Ricci-Tersenghi 2006] Achlioptas, D., Ricci-Tersenghi, F.: "On the solution space geometry of random constraint satisfaction problems"; Proceedings of the 36th ACM Symposium on Theory of Computing (2006), 130-139.

[Achlioptas et al. 2001a] Achlioptas, D., Chtcherba, A., Istrate, G., Moore, C.: "The phase transition in 1-in-k SAT and NAE3SAT"; Proceedings of the 12th ACM-SIAM Symposium on discrete algorithms (2001).

[Achlioptas et al. 2001b] Achlioptas, D., Kirousis, L.M., Kranakis, E., Krizanc, D.: "Rigorous results for random 2+p-SAT"; Theoretical Computer Science 265 (2001), 109-129.

[Achlioptas et al. 2004] Achlioptas, D., Beame, P., Molloy, M.: "A sharp threshold in proof complexity yields a lower bound for satisfiability search"; Journal of Computer and System Sciences 68, 2 (2004), 238-268.

[Bollobás 1985] Bollobás, B.: "Random Graphs"; Academic Press (1985).

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

[Braunstein et al. 2005] Braunstein, A., Mézard, M., Zecchina, R.: "Survey propagation: an algorithm for satisfiability": Random Structures and Algorithms 27, 2 (2005), 201-226.

[Creignou and Daudé 2004] Creignou, N., Daudé, H.: "Coarse and sharp thresholds for random generalized satisfiability problems"; Mathematics and Computer Science III: Algorithms, Trees, Combinatorics and Probabilities (Drmota, M., Flajolet, P., Gardy, D., Gittenberger, B., Editors), Birkhauser (2004), 507-517.

[Friedgut 1999] Friedgut, E.: "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 (1999) 1017-1054.

[Friedgut 2005] Friedgut, E.: "Hunting for sharp thresholds"; Random Structures and Algorithms 26, 1-2 (2005), 34-51.

[Gopalan et al. 2006] Gopalan, P., Maneva, E., Kolaitis, P., Papadimitriou, C.: "Computing the connectivity properties of the satisfiability solution space"; Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming, Lecture Notes in Computer Science. Springer Verlag (2006).

[Istrate 2005] Istrate, G.: "Threshold properties of random boolean constraint satisfaction problems"; Discrete Applied Mathematics 153 (2005), 141-152.

[Istrate et al. 2005] Istrate, G., Percus, A., Boettcher, S.: "Spines of random constraint satisfaction problems: Definition and connection with computational complexity"; Annals of Mathematics and Artificial Intelligence 44, 4 (2005), 353-372.

[Krzakala et al. 2007] Krzakala, F., Montanari, A., Ricci-Tersenghi, F., Semerjian, G., Zdeborova, L.: "Gibbs states and the set of solutions of random constraint satisfaction problems"; Proceedings of the National Academy of Sciences 104 (2007), 10318-10323.

[Maneva and Sinclair 2007] Maneva, E., Sinclair, A.: "On the satisfiability threshold and survey propagation for random 3-SAT"; Unpublished manuscript, available from first author's webpage (August 2007).

[Maneva et al. 2007] Maneva, E., Meltzer, T., Raymond, J., Sportiello, A., Zdeborova, L.: "A hike in the phases of the 1-in-3 satisfiability problem"; Technical Report cond-mat/0702421, arXiv.org (2007).

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

[Mézard et al. 2005a] Mézard, M., Mora, T., Zecchina, R.: "Clustering of solutions in the random satisfiability problem"; Physical Review Letters 94 (2005), 197205.

[Mézard et al. 2005b] Mézard, M., Mora, T., Zecchina, R.: "Pairs of SAT assignments and clustering in random boolean formulae"; Technical report, cond-mat/0506053, arXiv.org (2005).

[Mertens et al. 2006] Mertens, S., Mézard, M., Zecchina, R.: "Threshold values of random k-SAT from the cavity method" Random Structures and Algorithms 28, 3 (2006), 340-373.

[Molloy 2002] Molloy, M.: "Models for random constraint satisfaction problems"; Proceedings of the 32nd ACM Symposium on Theory of Computing (2002).

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

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

[Monasson et al. 1999b] Monasson, R., Zecchina, R., Kirkpatrick, S., Selman, B., Troyansky, L.: "Determining computational complexity from characteristic phase transitions" Nature 400, 8 (1999), 133-137.

[Raymond 2007] Raymond, J., Sportiello, A., Zdeborova, L.: "The phase diagram of random 1-in-3 satisfiability"; Phys. Rev. E 76 (2007), 011101.

[Rozenthal et al. 1999] Rozenthal, J., Plotkin, J.M., Franco, J.: "The probability of pure literals"; Journal of Logic and Computation 9, 4 (1999), 501-513.