Models and thresholds for random Constraint Satisfaction Problems.

Michael Molloy
Dept. of Computer Science
University of Toronto
Toronto, ON, Canada
molloy@cs.toronto.edu

November 12, 2001

Abstract

We introduce a class of models for random Constraint Satisfaction Problems. This class includes and generalizes many previously studied models. We characterize those models from our class which exhibit thresholds for satisfiability in the sense that the limiting probability of satisfiability changes significantly as the number of constraints increases. We also discuss models which exhibit sharp thresholds in the sense that the limiting probability jumps from 0 to 1 suddenly.

1 Introduction

A Constraint Satisfaction Problem (CSP) is a generalized form of Satisfiability which is widely studied in the Artificial Intelligence community. For example, the journal *Constraints* is devoted to these problems. Roughly speaking, a CSP generalizes SAT in the sense that variables can draw their values from a more general domain than simply $\{T, F\}$, and each clause (a.k.a. constraint) forms a set of restrictions as to which values the variables in that clause may jointly take.

It has long been observed that for $k \geq 3$, random instances of k-SAT drawn from a distribution close to the conjectured location of the threshold for satisfiability seem to be particularly hard to solve and so make very good test inputs for heuristics. (See [18] for an early reference; see [10, 3] for some theoretical work which supports this observation). Inspired by this fact, there has been a great deal of work in the AI community to generate hard CSP problems in a similar manner. In particular, there has been a lot of study towards estimating the satisfiability threshold for some natural models of random CSP (by generating instances with varying numbers of clauses and testing to see which ones tend to be satisfiable), and in studying the properties of random CSP's drawn from the region near a threshold. (See [19, 15, 17] as a few of many examples; see [14] for a more thorough list of 40+ papers.)

In [6] we proved a rather surprising fact: that most of the commonly studied models in fact had no thresholds of satisfiability! The appearance of such thresholds was just an illusion created by the fact that the instances tested all had a very small number of variables. (We describe this in more detail later in this introduction.)

This made it clear that there was a need for a mathematical analysis of the nature of thresholds in random models of CSP, and in particular a determination of which models exhibit satisfiability thresholds. That is the contribution of this paper.

Here, we introduce a general model for generating random Constraint Satisfaction Problems. We will generate problems on n variables, where each variable has a domain of $d \geq 2$ permissable values, and each constraint is on k variables for some fixed $k \geq 2$. For the sake of simplicity, we assume that each variable has the same domain, $\mathcal{D} = \{\delta_1, ..., \delta_k\}$, although this model can be easily modified to deal with more general situations. All asymptotics will be as $n \to \infty$, and so d, k are considered to be constants. Thus, typically n will be much larger than either d or k.

This model generalizes many previously studied models, including virtually all well-studied models of Constraint Satisfaction Problems, random instances of k-SAT, and CSP's which are equivalent to determining whether a random graph is k-colourable. Using straightforward methods to simulate constraints of size r by constraints of size t for any $t > r$, this model also generalizes models with mixed constraint sizes, such as the $(2+p)$-SAT model studied in [5].

1.1 Earlier models

We begin this subsection by more formally defining a CSP.

Given a k-tuple of variables, $(x_1, ..., x_k)$, a *restriction* on $(x_1, ..., x_k)$, is a k-tuple of values $R = (\delta_1, ...\delta_k)$ where each $d_i \in \mathcal{D}$. For each k-tuple $(x_1, ..., x_k)$, the set of restrictions on that k-tuple is called a *constraint*. The *empty constraint* is the constraint which contains no restrictions. We say that an assignment of values to the variables of a constraint C *satisfies* C if that assignment is not one of the restrictions in C. A CSP consists of a set of variables, a domain for each variable, and a set of constraints. An assignment of domain-values to all variables in a CSP satisfies that CSP if every constraint is simultaneously satisfied. The *constraint hypergraph* of a CSP is the k-uniform hypergraph whose vertices correspond to the variables, and whose hyperedges correspond to the k-tuples of variables which have *constraints*.

Until recently, the most commonly used random CSP models were the following:

Model A1: Specify M, p (typically $M = cn$ for some constant c and $0 < p \leq 1$ is independent of n). First choose a random k-uniform hypergraph on n vertices with M hyperedges (where each such hypergraph is equally likely), which will be the constraint hypergraph of our problem. Next, for each hyperedge e, we choose a constraint on the k variables of e as follows: Each of the d^k possible restrictions is chosen to be present with probability p (where each choice is, of course, independent of the corresponding choices for other potential contraints).

Model A2: Specify M, m (typically $M = cn$ for some constant c and $0 < p \leq 1$ is independent of n). Choose the constraint hypergraph as in Model A1. For each hyperedge e, we choose a constraint with exactly m restrictions uniformly at random from the set of $\binom{d^k}{m}$ such constraints.

Remark 1: Alternatively, we could have chosen the constraint hypergraphs by making an independent choice for each potential hyperedge, deciding to put it in the hypergraph with probability $p = \frac{c \times k!}{n^{k-1}}$. This variation produces models which are equivalent to Models A1 and A2 in the sense that the models will be a.s. satisfiable and a.s. unsatisfiable for the same values of c. A similar remark applies to the other two models described below.

In [6], we observed that Model A1 is asymptotically uninteresting in the sense that as long as M grows with n, the random CSP will be almost surely[1] (a.s.) unsatisfiable. The problem is quite simple: a.s. there will be at least one constraint which is overconstraining in that there is no assignment to its variables which satisfies even that constraint, i.e. a constraint which includes every possible restriction amongst its k variables. Model A2 avoids this particular problem, but has other equally damaging problems which implied that it is a.s. unsatisfiable so long as M is of order $\Theta(n)$ (as is usually the case) for all but small values of m: $m < d^{k-1}$ (Gent et al [14] showed that for these small values of m, the model does indeed exhibit a threshold).

Gent et al [14] provided a survey of 40+ recent AI papers which estimated the location of a satisfiability threshold for various models of random CSP, and/or studied random CSP's drawn from near the estimated threshold. They found that roughly 3/4 of papers used models which were shown in [6] to have no such threshold. Furthermore, those models which did have thresholds (mostly Model A2 with $m < d^{k-1}$) seem to have been "accidental" in the sense that the authors were unaware that the small value they chose for m was crucial.

In [6] we proposed an alternative model:

Model B: Specify M (typically $M = cn$ for some constant c). Choose a random CSP containing M restrictions, where each such CSP is equally likely to be chosen.

We showed that this model was asymptotically interesting in the sense that for small values of c a random CSP is a.s. satisfiable, while for large values

[1] A random CSP almost surely satisfies property P if $\lim_{n\to\infty} \mathbf{Pr}(P) = 1$.

of c it is a.s. unsatisfiable.

Nevertheless, this model is rather unsatisfactory in the sense that the CSP's which it generates do not resemble CSP's that researchers tend to find interesting. In particular, almost every constraint contains only one restriction, while in typical CSP's most constraints tend to have several restrictions.

Here, we propose a model which is both asymptotically interesting, and which overcomes the main disadvantage of Model B, in that it can generate CSP's whose constraints have several restrictions. We aim to keep the model as simple as possible, while at the same time meeting these two (somewhat conflicting) goals. In particular, our model is symmetric in the variables. It is not neccessarily symmetric amongst the values of $\mathcal{D}$, although the parameters of the model can be set so that it is.

Gent et. al. [14] proposed another model which is both asymptotically interesting and which generates reasonably complex constraints. Their model turns out to be a special case of the model introduced here. In the terminology introduced subsequently in this paper, they proved that their model is *very well behaved*, en route to showing that it exhibits a *transition*.

1.2 The new model

Our new model is similar in flavour to Models A1 and A2, but we must take some care in the way that we choose the constraints in order to avoid overconstrained constraints and similar problems. In particular, we cannot just independently choose whether or not to put each potential restriction in the constraint - we must choose the entire constraint en masse.

So, we consider a canonical set of variables, $X_1, ..., X_k$, and the set of $2^{d^k}-1$ potential non-empty constraints on $X_1, ..., X_k$. We specify a probability distribution $\mathcal{P}$ which selects a single random constraint. For simplicity, we assume that $\mathcal{P}$ is independent of n.

New Model: Specify $M, \mathcal{P}$ (typically $M = cn$ for some constant c). First choose a random constraint hypergraph with M hyperedges, in the usual manner. Next, for each hyperedge e, we choose a constraint on the k variables of e as follows: we take a random permutation from the k variables onto $\{X_1, ..., X_k\}$ and then we select a random constraint according to $\mathcal{P}$, mapping it onto a constraint on our k variables in the obvious manner.

Remark 2: Again, we could have chosen the constraint hypergraph by making an independent choice for each potential hyperedge, deciding to put it in the hypergraph with probability $p = \frac{c \times k!}{n^{k-1}}$. This variation produces a model which is, for our purposes, equivalent to the model as described above.

We set $\mathcal{C} = \text{supp}(\mathcal{P})$, i.e. the set of constraints, C, for which $\mathcal{P}(C) > 0$. As we shall see, we have the following interesting fact:

It turns out that whether or not this model exhibits a threshold depends only on $\mathcal{C}$. The actual distribution of $\mathcal{P}$ over $\mathcal{C}$ does not matter.

In what follows, we will specify the exact conditions that $\mathcal{C}$ must satisfy in order for the model to exhibit a threshold. To do so, we need a few definitions, starting with a more precise notion of a threshold.

We use $CSP_{n,M}(\mathcal{P})$, to denote a random CSP drawn from our New Model with parameters $n, M, \mathcal{P}$. We occsionally omit the subscript n, M, depending on the context. We say that $CSP(\mathcal{P})$ exhibits a *partial transition* if there exists constants $c_1, c_2 > 0$ such that

(i) if $M < c_1 n$ then $CSP_{n,M}(\mathcal{P})$ is not a.s. unsatisfiable;

(ii) if $M > c_2 n$, then $CSP_{n,M}(\mathcal{P})$ is a.s. unsatisfiable.

We say that $CSP(\mathcal{P})$ exhibits a *transition* if there exists $c_1, c_2 > 0$ such that

(i) if $M < c_1 n$ then $CSP_{n,M}(\mathcal{P})$ is a.s. satisfiable;

(ii) if $M > c_2 n$ then $CSP_{n,M}(\mathcal{P})$ is a.s. unsatisfiable.

Note that if $CSP(\mathcal{P})$ exhibits a transition then it exhibits a partial transition.

Remark: It turns out that there is no choice of $\mathcal{P}$ and c_2 such that for all $M > c_2 n$, $CSP_{n,M}(\mathcal{P})$ is neither a.s. satisfiable nor a.s. unsatisfiable. Thus there is no need for the obvious 3rd and 4th notions of "transition".

In the next two sections, we will characterize those models $CSP(\mathcal{P})$ which have transitions and partial transitions. We will see in that for each $\mathcal{C}$, we have that $CSP(\mathcal{P})$ exhibits a transition (partial transition) for at least one choice of $\mathcal{P}$ with $\mathcal{C} = \text{supp}(\mathcal{P})$ iff $CSP(\mathcal{P})$ exhibits a transition (partial transition) for *every* such choice of $\mathcal{P}$. Thus, it makes sense to say that $\mathcal{C}$ exhibits a *transition* or a *partial transition.*

In the fourth section, we will discuss models which have a sharp threshold for satisfiability.

A constraint C is *unsatisfiable* if it has no satisfying assignments, i.e. if it contains every possible restriction. A constraint C *forbids* $X_i = \delta$ if it has no satisfying assignments for which $X_i = \delta$. $X_i = \delta$ *implies* $X_j = \gamma$ *in* C if there are no satisfying assignments for C in which $X_i = \delta$ and $X_j \neq \gamma$. A constraint C *implies* a property P if every satisfying truth assignment for C meets P.

An *automorphism* of $\mathcal{C}$ is a bijection $\phi : \mathcal{D} \to \mathcal{D}$ such that $C = \{(\delta_1^1, ...\delta_k^1), ..., (\delta_1^t, ..., \delta_k^t)\} \in \mathcal{C}$ iff $\phi(C) = \{(\phi(\delta_1^1), ..., \phi(\delta_k^1)), ..., (\phi(\delta_1^t), ..., \phi(\delta_k^t))\} \in \mathcal{C}$. We say that $\mathcal{C}$ is *symmetric with respect to* $\mathcal{D}$ if it is symmetric in in the following sense: For any two values $\delta, \gamma \in \mathcal{D}$, there is an automorphism ϕ of $\mathcal{C}$ such that $\phi(\delta) = \gamma$.

Note that we do not require the symmetry to extend to $\mathcal{P}$, i.e. it is possible that $\mathcal{P}(C) \neq \mathcal{P}(\phi(C))$. Furthermore, $\mathcal{C}$ does not need to be symmetric in $X_1, ..., X_k$. For example, it could be that every constraint $C \in \mathcal{C}$ implies that $X_1 = X_2 \neq X_3$, while at the same time $\mathcal{C}$ is symmetric with respect to $\mathcal{D}$. However, the fact that our model takes a random mapping of the variables from an edge of the constraint hypergraph onto $(X_1, ..., X_k)$ implies that the model is symmetric with respect to the variables.

Our new model generalizes several previously studied models of CSP and other problems which reduce to being random instances of CSP. For example, Model A1 is simply this model where $\mathcal{P}(C) = p^{|C|}(1-p)^{d^k - |C|}$. Model A2 is the case where $\mathcal{P}(C) = \binom{d^k}{m}^{-1}$ if $|C| = m$ and $\mathcal{P}(C) = 0$ otherwise. Model B, is virtually equivalent to the case where $\mathcal{P}(C) = \frac{1}{d^k}$ if $|C| = 1$, and $\mathcal{P}(C) = 0$ otherwise. The well-studied problem of whether a random instance of k-SAT is satisfiable is equivalent to Model B where $d = 2$. The problem of deciding whether a random graph is c-colourable is equivalent to the case where $k = 2, d = c$ and $\mathcal{P}(C) = 1$ where $C = \{(\delta_1, \delta_1), ..., (\delta_c, \delta_c)\}$. In each of these examples, $\mathcal{C}$ is symmetric. It is already known that the first of these examples does not exhibit even a partial transition, that the second one only exhibits a partial transition if $m < d^{k-1}$, the third exhibits a transition, and the other two exhibit transitions for $k, c > 2$ and partial transitions for $k, c = 2$.

It is important to note that by the nature of the asymptotics involved, this study is only relevant as the number of variables grows very high. So these asymptotic results might be completely out-of-sync with experimental results which, due to technological limitations, are usually performed on a small number of variables, often less than 50. In fact, it is precisely this difference between the behaviour of the model on small instances and its asymptotic behaviour which led to earlier false assumptions about the asymptotic behaviour of Models A1 and A2.

2 The Symmetric Case

In this section, we focus on the case where $\mathcal{C}$ is symmetric. We say that a symmetric $\mathcal{C}$ is *well-behaved* if it satisfies the following two properties:

1. There is no $C \in \mathcal{C}, \delta \in \mathcal{D}$ and canonical variable X_i such that C forbids $X_i = \delta$.

2. For every $\delta \in \mathcal{D}$, there is at least one $C \in \mathcal{C}$ such that $X_1 = \delta, X_2 = \delta, ..., X_k = \delta$ does not satisfy C.

A symmetric $\mathcal{C}$ is *very well-behaved* if, in addition to properties 1,2, it also satisfies:

3. If a CSP is formed using constraints from $\mathcal{C}$, and its constraint hypergraph is a cycle, then it must be satisfiable.

It is worth noting that if Property 1 holds, then there is no unsatisfiable constraint $C \in \mathcal{C}$, and so we avoid the main problem with Model A1. However, as we will see, for $CSP(\mathcal{P})$ to not be a.s. unsatisfiable, it does not suffice to merely have no unsatisfiable constraint $C \in \mathcal{C}$ – we need the full power of Property 1 in order to avoid a more subtle problem.

Theorem 1 *If $\mathcal{C}$ is well-behaved, then for every distribution $\mathcal{P}$ with $\mathcal{C} = \text{supp}(\mathcal{P})$, $CSP(\mathcal{P})$ exhibits a* *partial transition. If $\mathcal{C}$ is not well-behaved, then for every distribution $\mathcal{P}$ with $\mathcal{C} = \mathrm{supp}(\mathcal{P})$, $CSP(\mathcal{P})$ does not exhibit a partial transition.*

Note that Theorem 1 immediately implies that Model A2 exhibits a partial transition iff $1 \le m < k^{d-1}$.

Theorem 2 *If $\mathcal{C}$ is very well-behaved, then for every possible distribution $\mathcal{P}$, $CSP(\mathcal{P})$ exhibits a transition. If $\mathcal{C}$ is not very well-behaved, then for every possible distribution $\mathcal{P}$, $CSP(\mathcal{P})$ does not exhibit a transition.*

Proof of Theorem 1: First we prove that our two properties are neccessary for $CSP(\mathcal{P})$ to have a partial transition.

Suppose that Property 2 does not hold. Then for at least one value δ, setting each $X_i = \delta$ will satisfy any $C \in \mathcal{C}$, and so setting every variable in our CSP equal to δ will satisfy our formula with probability 1, regardless of M (in fact, since $\mathcal{C}$ is symmetric in $\mathcal{D}$, this is true for every $\delta \in \mathcal{D}$). Thus, there is no value of M for which $CSP_{n,M}(\mathcal{P})$ is a.s. unsatisfiable.

Suppose that Property 1 does not hold, i.e. that there is some value $\delta \in \mathcal{D}$, canonical variable X_i, and $C \in \mathcal{C}$, such that C forbids $X_i = \delta$. Since $\mathcal{C}$ is symmetric in $\mathcal{D}$, we have that for *every* $\delta \in \mathcal{D}$ there is some $C(\delta) \in \mathcal{C}$ such that C forbids $X_i = \delta$.

We say that a variable x is *impossible* if x lies in at least d hyperedges of our constraint hypergraph, $e_1^j, ..., e_d^j$, and each e_ℓ^j has the constraint $C(\delta_\ell)$ where x is mapped onto X_i.

It is not hard to see that if x is impossible then our formula is not satisfiable, since x is forbidden from being assigned any of $\delta_1, ..., \delta_d$.

Suppose that $M = cn$ for any value of c.

Claim: A.s. $CSP_{n,M}(\mathcal{P})$ has an impossible variable.

Proof: First choose the constraint hypergraph. It is well known[2] (and an easy application of Chebychev's Inequality) that a.s. there will be at least ζn vertices with degree exactly d, such that no vertex lies in a 2-cycle, for a particular constant $\zeta = \zeta(c, k, d) > 0$. It is straightforward to show that a.s. at least one of these vertices will correspond to an impossible variable. We omit the easy details for the sake of space.

Thus, we have shown that if $\mathcal{C}$ is not well-behaved then either $CSP_{n,M}(\mathcal{P})$ is always a.s. unsatisfiable or it is always a.s. satisfiable, regardless of $\mathcal{P}$ and so $\mathcal{C}$ does not exhibit a partial transition.

Now we show that if $\mathcal{C}$ is well-behaved then it exhibits a partial transition for any choice of $\mathcal{P}$. First, we prove that Property 1 implies that we can take $c_1 = \frac{1}{k(k-1)}$:

It is well known[3] that for $c < \frac{1}{k(k-1)}$, the probability that our constraint hypergraph is acyclic tends to a constant $\rho(c) > 0$. It is not hard to show that if our constraint hypergraph is acyclic then Property 1 implies that our CSP is satisfiable. Thus the probability of satisfiability is at least ρ, and so we are not a.s. unsatisfiable.

For each component of the constraint hypergraph, we pick a variable to be the root, and set it arbitrarily. We then process each hyperedge in sequence, i.e. assign values to the variables of that hyperedge, such that every time we process a hyperedge, exactly one of its variables has already been set. Since no $C \in \mathcal{C}$ forbids any values from any variables, we will be able to find an assignment for the other variables on that hyperedge without violating the corresponding constraint.

Finally, we complete our proof by showing that Property 2 implies the existence of c_2, using what is by now a standard first moment argument.

There are d^n truth assignments. Consider any one particular assignment $\mathcal{A}$. There is a value δ such that some set S of at least $\frac{n}{d}$ variables are assigned δ in $\mathcal{A}$. By property 2, there is some constraint $C^* \in \mathcal{C}$ such that setting each $X_i = \delta$ violates C^*. The probability that $\mathcal{A}$ is a satisfying assignment is at most the probability that none of the k-tuples of S is constrained by C^*. Thus, if $M = cn$ and A is the number of satisfying truth assignments, then

$$\begin{aligned}\mathbf{Exp}(A) &\le d^n \left(1 - \frac{\mathcal{P}(C^*)}{d^k}\right)^{cn} \\ &< \mathrm{e}^{n\times\left(\ln d - \frac{c\mathcal{P}(C^*)}{d^k}\right)} \\ &= o(1)\end{aligned}$$

[2] See [8] for a proof of the case $k = 2$. The proof for higher values of k is similar.

[3] See [8] for a proof of the case $k = 2$. See also the branching argument inspired by Karp [16] in [7]. The proof for higher values of k is similar.

for $c > c_2 = \frac{d^k \ln d}{\mathcal{P}(C^*)}$. Thus, a.s. $A = 0$ for such c. □

Remark: In fact, note that we have shown that our CNF is always a.s. unsatisfiable iff Property 1 does not hold, and it is always a.s. satisfiable iff Property 2 does not hold.

Proof of Theorem 2: First we show why property 3 is neccessary for our model to have a transition. Suppose that it doesn't hold, i.e. that there is some unsatisfiable CNF formula C consisting of constraints from $\mathcal{C}$ and whose constraint hypergraph forms a ℓ-cycle.

For any value of $c > 0$, it is well known[4] that the probability that a random hypergraph with cn edges has an ℓ-cycle tends to a constant $\zeta_1 = \zeta_1(c, \ell) > 0$. If our constraint hypergraph has an ℓ-cycle, then the probability of that ℓ-cycle being constrained by C is equal to some constant $\zeta_2 > 0$. Therefore, the probability of being unsatisfiable tends to at least $\zeta_1\zeta_2 > 0$ and so we are not a.s. satisfiable for any value of c.

On the other hand, if property 3 does hold, then for $c < \frac{1}{k(k-1)}$ our CNF is a.s. satisfiable. It is well-knon that a.s. the constraint hypergraph will have no component with more than 1 cycle. If this happens, then the variables of each tree component can be set as in the proof of Theorem 1. The variables of each unicyclic component can be set as follows: first set the variables of the cycle in a satisfactory manner (which will always be possible by Property 3), and then set the remaining variables in the same manner as in a tree component.

The remainder of the theorem follows as in the proof of Theorem 1. □

In our proof, we showed that if $\mathcal{C} = \text{supp}(\mathcal{P})$ exhibits a transition then $CSP(\mathcal{P})$ is a.s. satisfiable when M is below the threshold for the constraint hypergraph to have a giant component. Not surprisingly, for many choices of $\mathcal{P}$, the same is still true for higher values of M. For example:

Theorem 3 *Suppose that $\mathcal{C}$ is such that for every minimally unsatisfiable formula F whose constraints are drawn from C, the ratio of constraints to variables of F is at least $\frac{1}{k-1}(1+\epsilon)$ for some constant $\epsilon > 0$. Then there is a constant $\delta > 0$ such that for any $\mathcal{P}$ with $\text{supp}(\mathcal{P}) = \mathcal{C}$, $CSP(\mathcal{P})$ is a.s. satisfiable for $M \le \frac{1}{k(k-1)}(1+\delta)n$.*

[4] See [8] for a proof of the case $k = 2$. The proof for higher values of k is similar.

Note that for the hypothesis of Theorem 3 to hold, $\mathcal{C}$ must be very well-behaved since a cycle has a hyperedge/vertex ratio of exactly $\frac{1}{k-1}$, and a tree has an even smaller ratio.

The proof of Theorem 3 is a standard type of argument in random graph theory. For the sake of space, we move it to the appendix.

Theorem 3 applies to several common models. For example, a minimally non-r-colourable graph has edge/vertex ratio at least $r/2$ since such a graph has minimum degree at least r. Also, a minimally unsatisfiable instance of k-SAT has clause/variable ratio at least $2/k$ since it cannot contain any pure literals. It would be interesting to characterize the distributions $\mathcal{P}$ for which $CSP(\mathcal{P})$ is a.s. unsatisfiable with $M = cn$ for all $c > \frac{1}{k(k-1)}$.

3 The asymmetric case

Here we extend the definitions and theorems of the previous section, to the case where $\mathcal{C}$ is not neccessarily symmetric in $\mathcal{D}$.

We say that a value δ is *0-bad*, if there is some canonical variable X_i, and constraint $C \in \mathcal{C}$ such that C forbids $X_i = \delta$. We say that δ is *j-bad* if there is some canonical variable X_i, and constraint $C \in \mathcal{C}$ such that C implies that if $X_i = \delta$ then at least one other canonical variable must be assigned a j'-bad value for some $j' < j$. A value is *bad* if it is j-bad for some j. A value is *good* if it is not bad.

If $\mathcal{C}$ is symmetric then, of course, if one value is bad then all values are bad, and so the existence of a bad value would violate Property 1 from the previous section. However, to allow for the case that $\mathcal{C}$ is asymmetric, we must modify Property 1. We say that $\mathcal{C}$ is *well-behaved* if it satisfies the following two properties:

1'. There is at least one good value in $\mathcal{D}$.

2. For every $\delta \in \mathcal{D}$, there is at least one $C \in \mathcal{C}$ such that $X_1 = \delta, X_2 = \delta, ..., X_k = \delta$ does not satisfy C.

Remark: It is an easy exercise to show that Properties 1' and 2 imply that there are at least 2 good values in $\mathcal{D}$.

We also must modify Property 3 slightly. We say that $\mathcal{D}$ is *very well-behaved* if, in addition to properties 1',2, it also satisfies:

3'. If a CSP is formed using constraints from $\mathcal{C}$, and its constraint hypergraph is a cycle, then it must have a satisfying assignment where no variable is assigned a bad value.

We have the analagous results from the previous section:

Theorem 4 *If $\mathcal{C}$ is well-behaved, then for every possible distribution $\mathcal{P}$, $CSP(\mathcal{P})$ exhibits a partial transition. If $\mathcal{C}$ is not well-behaved, then for every possible distribution $\mathcal{P}$, $CSP(\mathcal{P})$ does not exhibit a partial transition.*

Theorem 5 *If $\mathcal{C}$ is very well-behaved, then for every possible distribution $\mathcal{P}$, $CSP(\mathcal{P})$ exhibits a transition. If $\mathcal{C}$ is not very well-behaved, then for every possible distribution $\mathcal{P}$, $CSP(\mathcal{P})$ does not exhibit a transition.*

For the sake of space, we omit the proof of Theorem 5 which follows almost exactly like the proof of Theorem 2, and we defer the proof of 4 to the appendix.

Theorem 3 also applies to the asymmetric case. Again, it would be interesting to characterize the distributions $\mathcal{P}$ for which $CSP(\mathcal{P})$ is a.s. unsatisfiable with $M = cn$ for all $c > \frac{1}{k(k-1)}$.

4 Sharp thresholds

If a model $CSP(\mathcal{P})$ exhibits a transition, then it is natural to ask if it satisfies the following stronger property:

There exists a constant $c > 0$ such that for any $\epsilon > 0$, if $M < (c-\epsilon)n$ then $CSP(\mathcal{P})$ is a.s. satisfiable and if $M > (c+\epsilon)n$ then $CSP(\mathcal{P})$ is a.s. satisfiable.

The current "state-of-the-art" in the analysis of sharp thresholds does not allow us to prove a property this strong even for CNF-satisfiability. Instead, we must weaken the property slightly as follows.

$CSP(\mathcal{P})$ is said to have a *sharp threshold* of satisfiabilty if there exists a function $c(n)$ bounded away from 0 such that for any $\epsilon > 0$, if $M < (c(n) - \epsilon)n$ then $CSP_{n,M}(\mathcal{P})$ is a.s. satisfiable and if $M > (c(n) + \epsilon)n$ then $CSP_{n,M}(\mathcal{P})$ is a.s. satisfiable.

Note that the only difference between this and the stronger property is that here we allow c to vary with n. All "natural" models which are known to have sharp thresholds are conjectured to also satisfy the stronger property where $c(n) = c$ for all n. However, we do not know how to prove this.

This notion of sharp thresholds was introduced (in a broader context than that in the preceding definition) by Friedgut[12] who proved that random instances of k-SAT exhibit a sharp threshold. In doing so, he proved a general theorem characterizing random graph properties which have sharp thresholds. To describe this theorem, we must introduce some definitions.

A graph property, P, is *monotonically increasing* if (i) P is invarient under graph automorphisms, and (ii) whenever H is a subgraph of G, and H has P, then G must have P. This definition extends in the obvious manner to properties of CSP's - condition (ii) becomes: if F_1, F_2 are CSP's where every contraint in F_1 is also contained in F_2, and if F_1 has P then F_2 must have P. The only property which concerns us in this paper is that of being unsatisfiable, which is clearly monotonically increasing.

Given two properties P_1, P_2, their *symmetric difference*, $P_1 \triangle P_2$ is the property of satisfying one but not both of P_1, P_2. Given a set of graphs $\mathcal{H} = \{H_1, ..., H_k\}$, we define $Q(\mathcal{H})$ to be the property of having a subgraph which is isomorphic to one of the members of $\mathcal{H}$.

$G_{n,M}$ is the random graph with n vertices and M edges, where each such graph is equally likely to be chosen.

We extend our definitions of *transition* and *sharp theshold* to the setting of random graphs in the obvious manner. A monotonically increasing graph property, P, exhibits a transition if there exists $c_1, c_2 > 0$ such that if $M < c_1 n$ then a.s. $G_{n,M}$ does not have P, and if $M > c_2 n$ then a.s. $G_{n,M}$ has P. A graph property P with a transition has a sharp threshold if there exists a function $c(n)$ bounded away from 0 such that for any $\epsilon > 0$, if $M < (c(n) - \epsilon)n$ then a.s. $G_{n,M}$ does not have P, and if $M > (c(n) + \epsilon)n$ then a.s. $G_{n,M}$ has P.

Friedgut's Theorem yields the following:

Friedgut's Theorem[12][5] *If a monotonically increasing graph property P with a transition does not have a sharp threshold then there exists some $c(n)$ such that for every $\delta > 0$ there is a fixed set of unicyclic graphs $\mathcal{H} = \{H_1, ..., H_k\}$ (this set depends only on δ and not n) such that, when $M = c(n) \times n$, the probability of $G_{n,M}$ having $P \triangle Q(\mathcal{H})$ is less than δ.*

In other words, when $M = c(n)$, P can be arbitrarily closely approximated in probability by the property of having a subgraph from a list of unicyclic graphs.

This theorem also extends to random hypergraphs and to random CNF-formulae. It is natural to expect that it might extend to the random CSP-models discussed here, i.e. that $CSP(\mathcal{P})$ has a sharp threshold of satisfiability iff unsatisfiability cannot be arbitrarily well approximated by containing a subproblem isomorphic to one of a list of problems whose constraint hypergraphs are unicyclic. If $CSP(\mathcal{P})$ exhibits a transition then all problems with unicyclic constraint hypergraphs whose constraints are drawn from $\mathcal{C}$ are satisfiable. Thus, it would be natural to further expect that $CSP(\mathcal{P})$ has a sharp threshold of satisfiability iff it exhibits a transition. However, this is not the case, as the following example shows:

Example 1: Suppose that $k = 2, d = 4$, and that $\mathcal{C}$ contains exactly two constraints. The first one, C_1, forbids (X_1, X_2) from receiving any pair of values from $\{(1,1), (2,2), (3,3), (4,4)\}$. The second, C_2, forbids any pair (x, y) where one member of the pair is from $\{1,2\}$ and the other member is from $\{3,4\}$. $\mathcal{P}(C_1) = \frac{1}{3}$ and $\mathcal{P}(C_2) = \frac{2}{3}$.

Note that it is easy to verify that $CSP(\mathcal{P})$ exhibits a transition.

Suppose that $M = cn$ where $\frac{3}{4} < c < \frac{3}{2}$. Let G_i be the constraint graph induced by constraints of type C_i. A.s. G_1 will have $\frac{1}{3}M + o(M) = (\frac{1}{2} - c_1)n + o(n)$ edges and G_2 will have $(\frac{1}{2} + c_2)n + o(n)$ edges for some constants $c_1, c_2 > 0$. Therefore, by the most fundamental theorem in random graph theory (see [11]), a.s. no component of G_1 will have size greater than $O(\log n)$, while G_2 will have a *giant component* of size $\alpha n + o(n)$ for a constant $\alpha = \alpha(c) > 0$, and every other component of G_2 will have size at most $O(\log n)$.

Let T be the vertices of the giant component of G_2. Define B to be the event that G_1 has an odd cycle whose vertices all lie in T. It is easy to see that if B holds then our CSP is unsatisfiable: the C_2 constraints imply that either all variables in T recieve 1 or 2, or they all receive 3 or 4. However, if a set of vertices belongs to an odd cycle in G_1 then amongst them there must be at least 3 values. Furthermore, it is not difficult to verify that the probability that B does not hold AND the CSP is unsatisfiable is o(1). Therefore, the probability of satisfiability is $1 - \mathbf{Pr}(B) + o(1)$. It is also easily computed that $\mathbf{Pr}(B)$ tends to a constant (in terms of c) which is strictly between 0 and 1, and so our CNF is neither a.s satisfiable nor a.s. unsatisfiable - G_1 and G_2 are very close to being two independent random graphs, so this is esssentially the probability that a random graph with αn vertices and $(\frac{1}{2} - c_1)n$ edges contains an odd cycle. Therefore this model does not have a sharp threshold despite the fact that it exhibits a transition.

Dimitris Achlioptas[2] and Ehud Friedgut[13] pointed out to the author that for this particular example, there is a set of problems of the type which Friedgut's Theorem guarantees for random graphs and instances of k-SAT when a property has a coarse threshold. I.e., there is a set of problems such that having a subproblem from this set is a good approximation of being unsatisfiable. For each i, let $\mathcal{H}_i$ denote the set of subproblems such that, when G_1, G_2 are defined as above,

(i) G_1 is a cycle of length $2r + 1$ for some positive integer r, with vertices $v_1, ..., v_{2r+1}$;

(ii) G_2 consists of r disjoint trees $T_1, ..., T_r$ each of size i;

(iii) for each i, v_i is in T_i, and these are the only vertices that G_1, G_2 have in common.

Let P be the property that our random problem is unsatisfiable. For any $\delta > 0$, we can choose i large enough that $\mathbf{Pr}(P \triangle Q(\mathcal{H})) < \delta$. Almost surely, if

[5] Friedgut's Theorem is in fact much more general than this, in that (i) it shows that this is essentially an "if and only if" statement, and (ii) it allows for the possibility of a sharp threshold which occurs when M is not linear in n. We only include this version here for the sake of simplicity, and because it covers all situations which are relevant to this paper.

G_1 has an odd cycle lying in the giant component of G_2 then breadthfirst searchs, in G_2, from each vertex of the cycle will yield a collection of disjoint trees of length i, for any constant i. Therefore, $\mathbf{Pr}(P - Q(\mathcal{H}_i)) = o(1)$. On the other hand, the probability that G_1 has a cycle through the vertex of any non-giant component of G_2 which has size at least i, tends to a constant δ_i which tends to 0 as i grows. Choosing i so that $\delta_i < \delta$ yields $\mathbf{Pr}(Q(\mathcal{H}_i)) < \delta$. We omit the details, which are straightforward to a reader who is experienced in random graph theory.

It would be quite interesting to determine some neccessary and sufficient conditions on $\mathcal{P}$ for $CSP(\mathcal{P})$ to have a sharp threshold. We close this section by noting the following sufficient condition:

Theorem 6 *If*

(a) $\mathcal{C}$ is very well-behaved;

(b) every constraint in $\mathcal{C}$ has the property that for each value δ and canonical variable X_i, there is at most one restriction with $X_i = \delta$; and

(c) for each $C \in \mathcal{C}$, $\mathcal{P}(C) = \frac{1}{|\mathcal{C}|}$;

then $CSP(\mathcal{P})$ has a sharp threshold.

Remark 3: It is not difficult to weaken conditions (b) and (c) somewhat, but (a) is, of course, neccessary.

Note that this theorem generalizes the fact that random instances of k-SAT have a sharp threshold for $k \geq 3$ and the fact that k-colourability of random graphs has a sharp threshold for $k \geq 3$. This latter fact was proven by Achlioptas and Friedgut[4], and the proof of Theorem 6 is nearly identical to their proof. To avoid a long repetition of a proof which is readily available, we only give an outline of our proof.

Proof of Theorem 6: Rather than using Friedgut's Theorem, we we will make use of a similar theorem proved by Bourgain, subsequently to Friedgut's proof. The statement of this theorem is somewhat weaker, but it is easily applied to a more general situation, including the random models considered here.

Consider a set $\mathcal{A}$ of items $\{A_1, ..., A_r\}$, and a *selection probability* p. We will choose a random subset $A \subseteq \mathcal{A}$ as follows: for each A_i, we make an independent choice as to whether to include A_i in A, placing it in A with probability p. For example, if $\mathcal{A}$ is the set of possible edges of an n vertex graph, then we are simply choosing the random graph $G_{n,p}$, which is well known (see, for example, [8]) to be in many ways equivalent to $G_{n,M}$ where $M = p\binom{n}{2}$. Under the assumption that $\mathbf{Pr}(C)$ is rational for each $C \in \mathcal{C}$, we can also simulate $CSP(\mathcal{P})$ using this model as follows:

Following Remark 2, we will work in the equivalent model whereby we select each possible hyperedge to be in the constraint hypergraph independently with probability $p = \frac{c \times k!}{n^{k-1}}$.

$\mathcal{A}$ is the set of all $\binom{n}{k} k! |\mathcal{C}|$ possible constraints. We place each constraint into A with selection probability p' (to be named later), and we let our CSP consist of each constraint C such that (i) $C \in A$, and (ii) none of the other $k!|\mathcal{C}| - 1$ possible constraints on the same k-tuple of variables is in A. For any k-tuple of variables, that k-tuple forms a constraint with probability exactly $k!|\mathcal{C}| \times p'(1-p')^{k!|\mathcal{C}|-1}$, and we choose p' so that this value is equal to p. Furthermore, each member of $\mathcal{C}$ is equaly likely to be that constraint. Thus, this represents our model precisely. (Note that p' is very close to p - in fact $p' = p(1 + O(n^{-(k-1)}))$.)

This transformation allows us to apply Bourgain's Theorem to $CSP(\mathcal{P})$. As with Friedgut's Theorem, we will not state Bourgain's Theorem in it's full power - instead, we will only state it's implications to this setting, thus avoiding some technicalities.

Bourgain's Theorem [12] *Suppose that for each $c \in \mathcal{C}$, $\mathcal{P}(C) = \frac{1}{|\mathcal{C}|}$. If $CSP(\mathcal{P})$ does not have a sharp threshold of satisfiability then there exists $c(n)$ and absolute constants T, $0 < \gamma < 1$ and $\beta > 0$ such that when $M = c(n) \times n$, the probability that $CSP(\mathcal{P})$ is satisfiable is $\gamma + o(1)$ and either*

(i) with probability at least β, $CSP(\mathcal{P})$ contains an unsatisfiable subproblem A_1 of size at most T; or

(ii) there exists a satisfiable subproblem A_1 of size at most T such that conditioning on $CSP(\mathcal{P})$ containing A_1 lowers the probability of satisfiability to below $\gamma - \beta$.

A straightforward and well known[6] random graph

[6] See [8] for a proof of the case $k = 2$. The proof for higher

argument implies that a.s. every subproblem with at most T constraints is unicyclic. If $\mathcal{C}$ is very well behaved, all such problems are satisfiable. Thus, (i) does not hold and so we only need to eliminate the possibility of (ii).

So consider a CSP A_1 as in (ii). Without loss of generality, we can assume that A_1 is on the variables $x_1, ..., x_t$ for some $t \leq T$. Consider any satisfying assignment $x_1 = \delta_1, ..., x_t = \delta_t$ for A_1. The probability that $CSP(\mathcal{P})$ is satisfiable, conditional on $CSP(\mathcal{P})$ containing A_1 is at most the probability that $CSP(\mathcal{P})$ has a satisfying assignment with $x_1 = \delta_1, ..., x_t = \delta_t$.

Let $\mathcal{D}$ be the set of constraints, other than those in A_1, which involve $x_1, ...x_t$. A straightforward argument shows that there is an absolute constant R such that with probability at least $\frac{1-\beta}{2}$ (a) $|\mathcal{D}| \leq R$, (b) no constraint in $\mathcal{D}$ involves two variables from $\{x_1, ..., x_t\}$ and (c) no variable outside of $\{x_1, ..., x_t\}$ lies in more than one constraint in $\mathcal{D}$.

Consider any constraint $D \in \mathcal{D}$, say on variables $x_i, x_{j_1}, ..., x_{j_{k-1}}$ where $1 \leq i \leq t$ and each other index is greater than t. Since each constraint in $\mathcal{C}$ contains at most one restriction involving $x_i = \delta_i$, insisting that $x_i = \delta_i$ at worst reduces to imposing a single restricition of size $k-1$ on $x_{j_1}, ..., x_{j_{k-1}}$. This is no more restrictive than forbidding a single value for each x_{j_l}. Thus, conditional on (a), (b) and (c) holding, the probability that $CSP(\mathcal{P})$ has a satisfying assignment with $x_1 = \delta_1, ..., x_t = \delta_t$, is bounded from below by the following experiment:

Fix some collection of values $\zeta_1, ..., \zeta_{R\times(k-1)}$. Choose a random CSP from $CSP(\mathcal{P})$ with $M = c(n)$. Pick $R \times (k-1)$ variables at random, and add the additional requirements that the ith variable cannot recieve ζ_i.

Since the probability that either (a), (b) or (c) fails is at most $\frac{\beta}{2}$, condition (ii) implies that these additional requirements do not decrease the probability of satisfiability by more than $\frac{\beta}{2}$. We will show that this is impossible. To do so requires two steps:

(1) We show that if adding these $R \times (k-1)$ additional requirements decreases the probability of satisfiability by at least $\frac{\beta}{2}$, then so would increasing M from $c(n) \times n$ to $c(n) \times n + R'$, for an absolute constant R' defined in terms of $R \times (k-1)$.

(2) Increasing M from $c(n) \times n$ to $c(n) \times n + R'$ will only decrease the probability of unsatisfiability by $o(1)$.

Step (2) is at least intuitively obvious - adding a relatively very small number of constraints should not have a significant effect on the probability of satisfiability. However, it takes a little work to prove it.

The proof of Step (1) goes as follows: Suppose that we were only adding one additional requirement - we pick a random variable and forbid it from receiving ζ_1. Let X_1 be the set of variables such that, before this requirement, the only satisfying assignments set all of X_1 equal to ζ_1. If this requirement causes our problem to turn from satisfiable to unsatisfiable then we must have chosen a random variable in X_1. Thus, in order for this to have a significant probability of happening, X_1 would have to contain at least αn variables for some constant $\alpha > 0$ - suppose that this is the case. Then if instead of adding this extra requirement, we added a single additional constraint, we would cause the CSP to become unsatisfiable with probability at least $\alpha^k \times \frac{1}{|\mathcal{C}|}$ which is a positive constant. The reason is that this is the probability that the entire k-tuple is from X_1 and that we choose the restriction which forbids them from all being assigned the value ζ_1. Thus adding an additional constraint has a serious affect on the probability of satisfiability.

Modifying this argument to the case where we have $R(k-1)$ additional requirements uses the same ideas, but is is technically a little complicated.

The proofs of Steps (1) and (2) are identical to the corresponding steps found in [4], [12] and [1]. We refer the reader to either of those for more details. □

values of k is similar.

References

[1] D. Achlioptas. *Threshold Phenomena in Random Graph Colouring and Satisfiability.* Ph.D. Thesis, Dept. of Computer Science, University of Toronto, 1999.

[2] D. Achlioptas. Personal communication.

[3] D. Achlioptas, P. Beame and M. Molloy. *A sharp threshold in proof complexity.* Proceedings of STOC 2001, 337 - 346.

[4] D. Aclioptas and E. Friedgut. *A threshold for k-colourability.* Random Structures and Algorithms **14** (1999), 63 - 70.

[5] D. Achlioptas, L. Kirousis, E. Kranakis, D. Krizanc. *Rigorous results for random $(2+p)$-SAT.* Proceedings of RALCOM '97 (1997), 14 - 23.

[6] D. Achlioptas, L. Kirousis, E. Kranakis, D. Krizanc, M. Molloy, and Y. Stamatiou. *Random constraint satisfaction: a more accurate picture.* Constraints **6**, 329 - 324 (2001). Conference version in Proceedings of CP 97, 107 - 120.

[7] N. Alon and J. Spencer, *The Probabilistic Method.* Wiley (1992).

[8] B. Bollobas. *Random Graphs.* Academic Press (1985).

[9] P. Cheeseman, B. Kanefsky and W.M. Taylor. *Where the really hard problems are.* Proceedings of the 12th IJCAI (1991), 331 – 337.

[10] V. Chvatal and E. Szemeredi. *Many hard examples for resolution.* Journal of the ACM **35** (1988) 759 - 768.

[11] P. Erdos and A. Renyi. *On the evolution of random graphs.* Publication of the Mathematical Institute of the Hungarian Academy of Sciences, **5** (1960), 17–61.

[12] E. Friedgut and an appendix by J. Bourgain. *Sharp thresholds of graph properties and the k-SAT problem.* J. Am. Math. Soc. **12** (1999), 1017 - 1054.

[13] E. Friedgut. Personal communication.

[14] I. Gent, E. MacIntyre, P. Prosser, B. Smith and T. Walsh. *Random constraint satisfaction: flaws and structure.* Constraints **6**, 345 - 372 (2001).

[15] S. Grant and B.M. Smith. *The phase transition behaviour of maintaining arc consistency.* Proceedings of ECAI-96, 175 - 179.

[16] R. Karp. *The transitive closure of a random digraph.* Random Structures and Algorithms **1** (1990), 73 - 93.

[17] A. Meisels, S.E. Shimony and G. Solotorevsky. *Bayes networks for estimating the number of solutions to a CSP.* Proceedings of AAAI-97 179 - 184.

[18] D. Mitchell, B. Selman and H. Levesque. *Hard and easy distributions of SAT problems* Proceedings of AAAI-92, 459 - 465.

[19] P. Prosser. *An empirical study of phase transitions in binary constraint satisfaction problems.* Artificial Intelligence **81** (1996), 81 – 109.

A Two Proofs

Proof of Theorem 3 We can assume that ϵ is arbitrarily small, say $\epsilon \leq \frac{1}{2}$. Suppose that $M = cn$ where $c = \frac{1}{k(k-1)}(1+\delta)$, and $\delta < \epsilon$ is a small constant to be specified later. We will show that the expected number of subgraphs of our random constraint hypergraph which have the required hyperedge/vertex ratio is $o(1)$. Thus, a.s. none exist and so the CSP must be satisfiable since any unsatisfiable formula contains a minimally unsatisfiable subformula.

We can assume that such a subgraph has no tree or unicyclic components, since deleting such a component will only increase the hyperedge/vertex ratio. Thus, the subgraph must lie within the giant component. The size of the giant component is well-studied[7] and is a.s. at most αn where α tends to zero with δ. Furthermore Thus, we can restrict our attention to subgraphs of size at most αn.

For any $0 < a \leq \alpha n$, we consider the expected number of subgraphs on a vertices with $b = \frac{1}{k-1}(1+\epsilon)a$ edges. There are $\binom{a}{2}$ choices for the vertices, and $\binom{M}{b}$ choices for the edges. Given such choices, the probability that all the b edges lie within the a vertices is at most $(\frac{a}{n})^{kb}$. Therefore, the expected number of such subgraphs is at most:

$$\binom{n}{a}\binom{M}{b}\left(\frac{a}{n}\right)^{kb}$$

[7] See, for example, [8] for the case $k=2$. The case $k>2$ is amenable to similar analysis, or to the simpler branching process analysis found in [7] and inspired by [16].

$$
\begin{aligned}
&\leq \left(\frac{\mathrm{e}n}{a}\right)^a \left(\frac{\mathrm{e}M}{b}\right)^b \left(\frac{a}{n}\right)^{kb} \\
&= \left(\frac{\mathrm{e}^{1+(1+\epsilon)/(k-1)}c^{(1+\epsilon)/(k-1)}(k-1)}{1+\epsilon}\left(\frac{a}{n}\right)^{\epsilon}\right)^a \\
&= \left(K\left(\frac{a}{n}\right)^{\epsilon}\right)^a,
\end{aligned}
$$

where $K = K(\epsilon)$ is a constant. For $a \leq n^{\epsilon/4}$, this expression is $(\Theta(n^{-\epsilon(1-\epsilon/4)}))^a < n^{-\epsilon/2}$. And for $a < \alpha n$, it is less than $(\frac{1}{2})^a$ where α is sufficiently small in terms of ϵ. Thus the expected number of subgraphs is less than

$$\sum_{a=1}^{n^{\epsilon/4}} n^{-\epsilon/2} + \sum_{a \geq n^{\epsilon/4}} (\frac{1}{2})^a = o(1),$$

for $\delta = \delta(\alpha)$ sufficiently small. □

Proof of Theorem 4: First we show that if condition 1' does not hold, then for every distribution $\mathcal{P}$ and value $c > 0$, $CSP(\mathcal{P})$ is a.s. unsatisfiable.

For each $j \geq 0$ we let b_j denote the number of j-bad values. Thus $\sum_{j \geq 0} b_j = d$. For each 0-bad value δ, we let $C(\delta)$ be the constraint and $X_{i(\delta)}$ be the canonical variable such that $C(\delta)$ forbids $X_{i(\delta)} = \delta$. Similarly, if $j \geq 1$ then for each j-bad value δ, we let $C(\delta)$ be the constraint and $X_{i(\delta)}$ be the canonical variable such that if $X_{i(\delta)} = \delta$ then at least one other canonical variable must take a j'-bad value for some $j' < j$.

We define a *j-bad tree* rooted at a variable v as follows:

- A 0-bad tree rooted at v is a set of b_0 hyperedges intersecting at v such that for each 0-bad value δ, one of the hyperedges receives the constraint $C(\delta)$ with $X_{i(\delta)}$ mapped onto v.
- For $j \geq 1$, a j-bad tree rooted at v consists of a j'-bad tree $T_{j'}$ rooted at v for each $0 \leq j' < j$, along with b_j hyperedges $e_1, ..., e_{b_j}$ intersecting at v such that:
  - (i) for each j-bad value δ, a hyperedge e_i receives the constraint $C(\delta)$ with $X_{i(\delta)}$ mapped onto v; and
  - (ii) for every value $u \neq v$ in one of the e_i, u is the root of a $(j-1)$-bad tree T_u; and
  - (iii) all of the bad trees $\{T_0, ..., T_{j-1}\} \cup \{T_u :$ u is in some $e_i\}$ are disjoint except that $T_0, ..., T_{j-1}$ all contain v.

It is not hard to argue inductively that if v roots a j-bad tree then v cannot recieve a j'-bad value for any $j' \leq j$: If $j = 0$ then for each 0-bad value δ, v is forbidden δ by $C(\delta)$. For $j \geq 1$, since for each $j' < j$, v roots a j'-bad tree, v cannot receive a j'-bad value. if v recieves some j-bad value, δ, then since v lies in a $C(\delta)$-constraint, at least one of the other variables in that constraint must receive a j'-bad value for some $j < j$. But each of those variables roots a j'-bad tree and thus cannot take any such bad value.

A *bad tree* is a j-bad tree where j is the maximum value such that $b_j > 0$. If 1' does not hold, then every value is a j'-bad value for some $j' < j$ and so the root of a bad tree cannot receive any of them. Therefore, if a CSP contains a bad tree then it is unsatisfiable.

Now, we show that for any distribution $\mathcal{P}$ and $c > 0$, a.s. $CSP(\mathcal{P})$ contains a bad tree. The following fact is well known in random graph theory.

Fact: For any fixed hypertree T independent of n, there exists a constant $z > 0$ such that a.s. our random constraint hypergraph will contain at least ζn copies of T.

The basic outline for the proof of this fact is as follows: (1) The expected number of copies of T is $\zeta' n$ by a simple but tedious calculation. (2) Because n grows much larger than the (constant) size of T, individual potential appearances of T occur nearly independently. An application of Chebychev's inequality formalizes this and shows that the number of copies of T is highly concentrated. In particular, it is a.s. at least ζn for any $\zeta < \zeta'$. See [8] for details of the case $k = 2$. The analysis for higher k is the same. For the sake of space, we omit the straightforward details.

The remainder of the proof, namely the relevance of property 2, is identical to the proof of Theorem 1. □