Sharp thresholds for constraint satisfaction problems and homomorphisms

Hamed Hatami *and* Michael Molloy
Department of Computer Science
University of Toronto
Toronto, Canada
{hamed,molloy}@cs.toronto.edu

November 4, 2004

Abstract

We determine under which conditions certain natural models of random constraint satisfaction problems have sharp thresholds of satisfiability. These models include graph and hypergraph homomorphism, the (d,k,t)-model, and binary constraint satisfaction problems with domain size 3.

1 Introduction

Random 3-SAT and its generalizations have been studied intensively for the past decade or so (see eg. [1, 5, 8, 17, 25, 2, 7, 13, 15, 33, 9]). One of the most interesting things about these models, and arguably the main reason that most people study them, is that many of them exhibit what is called a *sharp threshold of satisfiability*[1], a critical clause-density at which the random problem suddenly moves from being almost surely[1] satisfiable to almost surely unsatisfiable. Most of the work on these problems is, at least implicitly, an attempt to determine the precise locations of their thresholds. At this point, these locations are known only for a handful of the problems, such as [2, 7, 13, 15, 33, 9]. Just proving the existence of a sharp threshold for random 3-SAT was considered a major breakthrough by Friedgut[17]. The vast majority of these generalizations appear to have sharp thresholds, but there are exceptions which are said to have coarse thresholds[1].

The ultimate goal of the present line of enquiry is to determine precisely which of these models have sharp thresholds, but this appears to be quite difficult; in Section 2 we show that it is at least as difficult as determining the location of the threshold for 3-colourability, something that has been sought after for more than 50 years (see eg. [14, 28]). A more fundamental goal is to obtain a better understanding of what can cause some problems to have coarse thresholds rather than sharp ones.

Molloy[27] and independently Creignou and Daude[10] introduced a wide family of models for random constraint satisfaction problems which includes 3-SAT and many of its generalizations. This permits us to study them under a common umbrella, rather than one-at-a-time. Molloy determined precisely which models from this family have any threshold at all ([10] provides the same result for those models with domain[1] size 2). But he left open the much more important question of which models have sharp thresholds. In this paper, we begin to address this question. We answer it for two of the most natural subfamilies - the so-called (d,k,t)-family[1] (Theorem 2), and the family of graph and hypergraph homomorphism problems (Theorem 3). We also shed light on the more fundamental problem by determining the only properties that can cause a coarse threshold in binary constraint satisfaction problems with domain size 3.

The standard example of a problem with a coarse threshold is 2-colourability. Here, there is a coarse threshold precisely because unsatisfiability (i.e. non-2-colourability) can be caused only by the presence of odd cycles. Roughly speaking, Friedgut's Theorem[17] implies that a problem exhibits a coarse threshold iff unsatisfiability is *approximately* equivalent to having one of a set of unicyclic[1] subproblems. It is not hard to see that if there are unsatisfiable unicyclic instances of a problem then that problem exhibits a coarse threshold (or exhibits no threshold at all). This makes it quite natural to pose the following rule-of-thumb:

Hypothesis A: *If a random model from the family in [27] is such that: (a) it exhibits a threshold, and (b) every unicyclic instance is satisfiable, then that threshold is sharp.*

However, reality is not that simple. [27] presents a counterexample to Hypothesis A; others are presented in this paper. Nevertheless, the hypothesis holds for certain subfamilies of models. Creignou and Daude[10] conjectured that Hypothesis A holds for problems with domain-size two; this was proven by Istrate[23] and independently Creignou and Daude[11] proved it for the case where the model is *symmetric*. Theorems 2 and 3 in this paper show that Hypothesis A holds for the (d,k,t)-models and for homomorphism problems.

In general, coarse thresholds can be caused by much more subtle and insidious reasons than unsatisfiable unicyclic instances. In this paper we begin to understand some of these reasons by focusing

[1]Defined formally below.

on the case where the constraint size is two and the domain size is three (a natural next step after the well-understood domain-size-two case). In this paper, we identify a particular subtle cause, and show that this and unsatisfiable unicyclic instances are the only things that can cause a coarse threshold (Theorem 9). If we permit either greater domain sizes or greater constraint sizes then this is no longer true - there are other possible causes.

1.1 The random models

In our setting, the variables of a constraint satisfaction problem (CSP) all have the same domain of permissable values, $\{1, ..., d\}$, and all constraints will have size k, for some fixed integers d, k. 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 $1 \leq \delta_i \leq d$. 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. An assignment of values to all variables in a CSP *satisfies* that CSP if every constraint is simultaneously satisfied. A CSP is *satisfiable* if it has such a satisfying assignment.

It will be convenient to consider a set of canonical variables $X_1, ..., X_k$ which are used only to describe the "pattern" of a constraint. These canonical variables are not variables of the actual CSP. For any d, k there are d^k possible restrictions and 2^{d^k} possible constraints over the k canonical variables. We denote this set of constraints as $\mathcal{C}^{d,k}$. For our random model, one begins by specifying a particular probability distribution, $\mathcal{P}$ over $\mathcal{C}^{d,k}$. We use $\mathbf{supp}(\mathcal{P})$ to denote the support of $\mathcal{P}$; i.e. the set of constraints C with $\mathcal{P}(C) > 0$. Different choices of $\mathcal{P}$ give rise to different instances of the model.

We now define our random models. The "$G_{n,M}$" model, where the number of constraints is fixed to be M, is the most common. But in this paper, it will be much more convenient to focus on the "$G_{n,p}$" model where each k-tuple of variables is chosen independently with probability $p = c/n^{k-1}$ to receive a constraint. The two models are, in most respects, equivalent when $M = (c/k!)n$. In particular, it is straightforward to show that one exhibits a sharp threshold iff the other does.

The $CSP_{n,p}(\mathcal{P})$ Model: Specify n, p and $\mathcal{P}$ (typically $p = c/n^{k-1}$ for some constant c; note that $\mathcal{P}$ implicitly specifies d, k). First choose a random k-uniform hypergraph on n variables where each of the $\binom{n}{k}$ potential hyperedges is selected with probability p. 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 $\mathcal{P}$ and map it onto the k variables.

A property holds *almost surely* (a.s) if the limit as $n \to \infty$ of it holding is 1. We say that $CSP_{n,p}(\mathcal{P})$ has a *sharp threshold of satisfiability* if there is some $c = c(n) > 0$ such that for every $\epsilon > 0$, if $p = (1-\epsilon)c/n^{k-1}$ then $CSP_{n,p}(\mathcal{P})$ is a.s. satisfiable and if $p = (1+\epsilon)c/n^{k-1}$ then $CSP_{n,p}(\mathcal{P})$ is a.s. unsatisfiable. This is often abbreviated to just *sharp threshold*. We say that $CSP_{n,p}(\mathcal{P})$ has a *coarse threshold* if for all c in some interval $c_1(n) < c < c_2(n)$, $CSP_{n,p}(\mathcal{P})$ is neither a.s. satisfiable nor a.s. unsatisfiable. If $CSP_{n,p}(\mathcal{P})$ has neither a sharp nor a coarse threshold, then it is easy to see that it must either be a.s. satisfiable for all $c > 0$ or a.s. unsatisfiable for all $c > 0$.

Each k-tuple of vertices can have at most one constraint in $CSP_{n,p}(\mathcal{P})$. When applying Friedgut's theorem, it will be convenient to relax this condition, and allow k-tuples to possibly receive multiple constraints. Thus up to $k! \times |\mathbf{supp}(\mathcal{P})|$ constraints can appear on a k-tuple of variables.

The $\widehat{CSP}_{n,p}(\mathcal{P})$ Model: Specify n, p and $\mathcal{P}$. For each of the $n(n-1)...(n-k+1)$ *ordered* k-tuples of variables and each constraint $C \in \mathbf{supp}(\mathcal{P})$, we assign C to the ordered k-tuple with probability $\mathcal{P}(C) \times p/k!$.

Note that the expected total number of constraints is the same under each model. Furthermore, it is easy to calculate that the probability of at least one k-tuple receiving more than one constraint in $\widehat{CSP}_{n,p}(\mathcal{P})$ is for $k \geq 3$, $o(1)$ and for $k = 2$, an absolute constant $0 < \alpha < 1$. It follows that if a property holds a.s. in $\widehat{CSP}_{n,p}(\mathcal{P})$ then it holds a.s. in $CSP_{n,p}(\mathcal{P})$. As a corollary, we have:

Lemma 1 *If $\widehat{CSP}_{n,p}(\mathcal{P})$ has a sharp threshold then so does $CSP_{n,p}(\mathcal{P})$. The reverse is true for $k \geq 3$.*

So for the remainder of the paper, whenever we wish to prove that $CSP_{n,p}(\mathcal{P})$ has a sharp threshold, we will work in the $\widehat{CSP}_{n,p}(\mathcal{P})$ model.

We often focus on the *constraint hypergraph* of a CSP; i.e. the hypergraph whose vertices are the variables and whose edges are the tuples of variables that have constraints. A *tree-CSP* is a CSP whose constraint hypergraph is a hypertree. A CSP is *unicyclic* if its constraint hypergraph is unicylic; i.e. has exactly one cycle. (Hypertree and cycle are defined below).

We close this subsection with some hypergraph definitions. A hypergraph consists of a set of vertices and a set of *hyperedges*, where each hyperedge is a collection of vertices. If every hyperedge has size exactly k then the hypergraph is *k-uniform.* In a *simple* hypergraph, no vertex appears twice in any one hyperedge, and no two edges are identical. So, for example, the constraint hypergraph of $CSP_{n,p}(\mathcal{P})$ is simple, but the constraint hypergraph of $\widehat{CSP}_{n,p}(\mathcal{P})$ may have multiple edges. Neither model permits multiple copies of a vertex in a single edge, but such edges are possible when we discuss hypergraph homomorphism problems. The edge $(v, v, ..., v)$ is called a *loop.*

A walk P of size r is a sequence of r hyperedges and $r+1$ vertices $(v_0, e_1, v_1, e_2, v_2 \ldots, e_r, v_r)$ such that e_i contains both v_{i-1} and v_i. A walk is a *path* if the v_i are distinct. A walk is a *cycle* of size r if for $i = 1, \ldots, r$ the v_i are distinct, and $v_0 = v_r$. The *distance* from a vertex u to a vertex v is the minimum r such that there exists a walk of length r, $(v, e_1, v_1, \ldots, e_r, u)$; the distance of a vertex from itself is defined to be 0. The distance from a vertex v to a set of vertices is the minimum distance from v to any vertex in the set. A hypergraph is a *hypertree* if it has no cycles and it is connected.

By *contracting* two vertices u and v into a new vertex w, we mean (i) adding a new vertex w to the set of the vertices, (ii) replacing u and v in every hyperedge by w, and (iii) removing u and v.

1.2 Two special families

Perhaps the most natural choice for $\mathcal{P}$ is the distribution obtained by selecting each of the d^k possible restrictions independently with probability q for some constant q. However, as noted in [4], every such choice of $\mathcal{P}$ yields a model that is a.s. unsatisfiable for any non-trivial choice of p. So this is a rather uninteresting family of models, particularly as far as the study of thresholds goes.

The next most natural choice for $\mathcal{P}$ is to fix t, the number of restrictions per clause, and to make every constraint with exactly t restrictions equally likely. (Note that for $d = 2, t = 1$ this yields random k-SAT.) This is often called the (d, k, t)-model and has received a great deal of study, both from a theoretical perspective [26, 30] and from experimentalists (see [18] for a survey of many such studies). In [4] it is shown that when $t \geq d^{k-1}$, this model is problematic in the same way as the previously mentioned one, as it is a.s. unsatisfiable even for values of $p = o(1/n^{k-1})$ (i.e. when the number of constraints is $o(n)$). However, it was proven in [18] that for every $1 \leq t < d^{k-1}$, the (d, k, t)-model does not have that problem. One of the main contributions of this paper is to show that in this case the model exhibits a sharp threshold:

Theorem 2 *For every $d, k \geq 2$ and every $1 \leq t < d^{k-1}$, the (d, k, t)-model has a sharp threshold.*

From a different perspective, it is quite natural to consider the case where every constraint is identical, i.e. $|\mathbf{supp}(\mathcal{P})| = 1$. It is not hard to see that every such problem is equivalent to a hypergraph homomorphism problem, as defined below:

For two k-uniform hypergraphs, G, H, a *homomorphism* from G to H is a mapping h from $V(G)$ to $V(H)$ such that for each edge $(v_1, v_2, \ldots, v_k)$ of G, $(h(v_1), h(v_2), \ldots, h(v_k))$ is an edge of H. We say that G is *homomorphic* to H, if there exists such a homomorphism. When $k = 2$ and H is the complete graph with no loops, we are simply asking whether G has a d-colouring. Homomorphisms are an important generalization of graph colouring (see, eg [21]). They are often also referred to as H-colourings (eg [22, 20]).

Suppose that H is a fixed hypergraph, and G is a random hypergraph on n vertices where each of the $\binom{n}{k}$ potential hyperedges is selected with probability p. Set d to be equal to the number of vertices in H and define a constraint C with domain size d and constraint size k by saying that C permits $X_1 = \delta_1, ..., X_k = \delta_k$ iff $(d_1, ..., d_k)$ is a hyperedge of H. Treat each vertex of G as a variable with domain $\{1, .., d\}$ and assign C to each hyperedge of G. We call this the *H-homomorphism problem.*

Thus we have an instance of $CSP_{n,p}(\mathcal{P})$ where C is the only constraint in $\mathbf{supp}(\mathcal{P})$ and furthermore C is symmetric under permutations of the canonical variables; in other words, all constraints are identical even under permutations of variables. It is easy to see that every such $\mathcal{P}$ corresponds to a homomorphism problem; just take H to be the hypergraph where $(\delta_1, ..., \delta_k)$ is a hyperedge iff C permits $X_1 = \delta_1, ..., X_k = \delta_k$. Note that here a hyperedge may contain multiple copies of a vertex.

Thus, these H-homomorphism problems are not only important as a fundamental graph problem, but also because they form a very natural subclass of our family of random CSP models. In this paper, we prove that Hypothesis A holds for every connected H.

It is easy to see that if H has a loop $(\delta, \delta, ..., \delta)$ then every hypergraph is trivially homomorphic to H (just map every vertex to δ); so the H-homomorphism problem has no threshold at all. The other trivial case is where H has no hyperedges at all and so no non-trivial hypergraph has an H-homomorphism.

It is an easy exercise to show that: (1) For $k \geq 3$, every unicylic subgraph is homomorphic to a single hyperedge, and hence to H. (2) For $k = 2$: if the triangle is homomorphic to H, then so is every unicyclic graph; and the triangle is homomorphic to H iff H contains a triangle. Therefore, proving that Hypothesis A holds whenever H is connected and undirected is equivalent to proving:

Theorem 3 *If H is a connected undirected loopless hypergraph with at least one edge, then the H-homomorphism problem has a sharp threshold iff (a) $k \geq 3$ or (b) $k = 2$ and H contains a triangle.*

We do not have a strong feeling as to whether the "connected" condition is necessary here; we discuss the possibility of extending Theorem 3 to disconnected graphs in Section 3.

1.3 Tools

Our main tool is distilled from Friedgut's main theorem in [17]. Friedgut reported to us[19] that his proof can be adapted to the setting of this paper. That adaptation does not appear in print, so for completeness, we provide one in the appendix. There, we also describe how by applying that theorem to the setting of this paper, we obtain the following lemma. Roughly speaking, it says that for a value of p that is "within" the coarse threshold, there is a constant sized unicyclic instance M such that adding M to our random CSP boosts the probability of unsatisfiability by at least $2\alpha > 0$, whereas adding a linear number of new random constraints only boosts it by at most α.

First, we must formalize what we mean by "adding M". Given two CSP's M, F where M has r variables, and F has at least r variables, we define $F \oplus M$ to be the CSP obtained by choosing a random r-tuple of variables in F and then adding M on those r variables.

Lemma 4 *For any $\mathcal{P}$, if $\widehat{CSP}_{n,p}(\mathcal{P})$ has a coarse threshold of satisfiablity then there exist $p = p(n)$, $\alpha, \epsilon > 0$, and a unicyclic CSP M on a constant number of variables whose constraints are chosen from* $\mathbf{supp}(\mathcal{P})$ *such that:*

(a) $\alpha < \Pr[\widehat{CSP}_{n,p}(\mathcal{P})$ *is unsatisfiable*$] < 1 - 3\alpha$.

(b) $\Pr[\widehat{CSP}_{n,p(1+\epsilon)}(\mathcal{P})$ *is unsatisfiable*$] < 1 - 2\alpha$.

(c) $\Pr[\widehat{CSP}_{n,p}(\mathcal{P}) \oplus M$ *is unsatisfiable*$] > 1 - \alpha$.

Our second tool is easily proven with a straightforward first moment calculation and concentration argument (via eg. the second moment method or Talagrand's Inequality). We omit the proof.

Lemma 5 *Let T be a tree-CSP whose constraints are in* $\mathbf{supp}(\mathcal{P})$. *There exists $z = o(1/n)$ such that a.s. $\widehat{CSP}_{n,z}(\mathcal{P})$ contains T as a sub-CSP.*

2 Difficulty

The ultimate goal of this research is to characterize all distributions $\mathcal{P}$ for which $CSP_{n,p}(\mathcal{P})$ exhibits a sharp threshold. However, the following example indicates that this is very difficult, even for binary CSP's (the case where $k = 2$). In particular, it is at least as difficult as determining the location of the 3-colourability threshold, a heavily pursued open problem. (The existence of that threshold was proven in [3]; see [28] for a recent survey and see [6, 16] for the best current bounds on its location.)

We set $d = 5$ and $k = 2$, and define two constraints by listing their pairs of forbidden values:

$$\begin{aligned} C_1 &= \{(1,1),(2,2)\} \cup (\{1,2\} \times \{3,4,5\}) \cup (\{3,4,5\} \times \{1,2\}), \\ C_2 &= \{(3,3),(4,4),(5,5)\} \cup (\{1,2\} \times \{3,4,5\}) \cup (\{3,4,5\} \times \{1,2\}). \end{aligned}$$

Note that each constraint forces the endpoints of every edge to take values that are either both in $\{1,2\}$ or both in $\{3,4,5\}$. A C_1 constraint says that they have to be different values if they are both in $\{1,2\}$. A C_2 constraint says that they have to be different values if they are both in $\{3,4,5\}$.

We let C_1 occur with probability q and C_2 occur with probability $1-q$ in $\mathcal{P}$. Set $c(q) = (1-q)/q$.

Fact 6 *(a) If $G_{n,p=c(q)/n}$ is a.s. 3-colourable, then $CSP_{n,p}(\mathcal{P})$ has a sharp threshold.*

(b) If there is some $\epsilon > 0$ such that $G_{n,p=(c(q)-\epsilon)/n}$ is a.s. not 3-colourable, then $CSP_{n,p}(\mathcal{P})$ has a coarse threshold.

Thus, determining the type of the threshold for all such models $CSP_{n,p}(\mathcal{P})$ requires the knowledge of for which values of c, $G(n, \frac{c}{n})$ is a.s. 3-colourable, and for which values it is a.s. not 3-colourable.

Proof Choose our $CSP_{n,p}(\mathcal{P})$ by first taking $G_{n,p=c/n}$ and then setting each edge to be C_1 with probability q and C_2 otherwise. Let G_1, G_2 be the subgraphs formed by the edges chosen to be C_1, C_2

respectively. If $c < 1$ then all components of $G_{n,p=c/n}$ are trees or unicycles and the CSP is trivially satisfiable. So we can focus on the range $c > 1$ and we let T denote the giant component of $G_{n,p=c/n}$. Note that the variables of T must either all take values from $\{1,2\}$ or all take values from $\{3,4,5\}$.

Case 1: $c > \frac{1}{q}$. Then G_1 is equivalent to $G_{n,p=c_1/n}$ for some $c_1 > 1$ and it follows easily that a.s. G_1 contains a giant component which is not 2-colourable. This giant component is a subgraph of T and so the variables of T must all take values from $\{3,4,5\}$. If follows that the CSP is satisfiable iff G_2 is 3-colourable. Note that G_2 is equivalent to $G_{n,p=c_2/n}$ for some $c_2 > c(q)$.

Case 2: $c < \frac{1}{q}$. Then G_1 is equivalent to $G_{n,p=c_1/n}$ for some $c_1 < 1$ and G_2 is equivalent to $G_{n,p=c_2/n}$ for some $c_2 < c(q)$. If G_2 is a.s. 3-colourable then the CSP is a.s. satisfiable. If G_2 is a.s. not 3-colourable then the CSP is satisfiable iff T is 2-colourable; i.e., if G_2 does not have an odd cycle lying within T. It is easy to see that this occurs with probability between ζ and $1-\zeta$ for some $\zeta > 0$; i.e. that the CSP is neither a.s. satisfiable nor a.s. unsatisfiable.

Fact 6 now follows. If $G_{n,p=c(q)/n}$ is a.s. 3-colourable, then $CSP_{n,p}(\mathcal{P})$ has a sharp threshold which lies somewhere above $\frac{1}{q}$. If there is some $\epsilon > 0$ such that $G_{n,p=(c(q)-\epsilon)/n}$ is a.s. not 3-colourable, then $CSP_{n,p}(\mathcal{P})$ has a coarse threshold running from $\frac{1}{q} - \delta$ to $\frac{1}{q}$ for some $\delta > 0$. □

3 Homomorphisms

In this section, we prove our theorem concerning H-homomorphisms. Let $G^k_{n,p}$ denote the random k-uniform hypergraph on n vertices where each k-tuple is present as a hyperedge with probability p.

Proof of Theorem 3 We begin with the case $k \geq 3$. Let H be some k-uniform hypergraph, and assume that H has a coarse threshold. Let M, p, α, ϵ be as guaranteed by Lemma 4. In this setting, M is a unicyclic k-uniform hypergraph, such that adding M to $G^k_{n,p}$ boosts the probability of not having a homomorphism to H by at least 2α.

Consider $G = G^k_{n,p} \oplus M$. Let M^+ be the subgraph of G consisting of all hyperedges that contain at least one vertex of M (and, of course all vertices in those hyperedges); in other words, M^+ is the subhypergraph induced by the vertices of M and all their neighbours. Standard properties of random hypergraphs (or straightforward first moment calculations) imply that there is some constant L such that with probability at least $1 - \frac{\alpha}{2}$: M^+ is unicyclic and has at most L vertices.

Since M is unicylic, it is an easy exercise to show that there exists a homomorphism h from M to a single edge, say $(v_1, \ldots, v_k)$. Let h_i be the set of the vertices in M that are mapped by h to v_i. Obtain the hypergraph G' from G by (i) removing all edges in M; (ii) contracting all of the vertices in h_i into one single new vertex u_i, for each $1 \leq i \leq k$; (iii) adding the single hyperedge $(u_1, \ldots, u_k)$.

Suppose that h' is a homomorphism from G' to H. Then a mapping from the vertices of G to the vertices of H which maps every vertex v in $G - M$ to $h'(v)$, and every vertex in h_i to $h'(u_i)$ is a homomorphism from G to H. Thus, if G' is homomorphic to H then so is G.

Let T be the hypertree defined as follows: T has a hyperedge $t_1, ..., t_k$, and each t_i lies in L other hyperedges. Only $t_1, ..., t_k$ lie in more than one edge of T. Thus, T has $k + k(k-1)L$ vertices and $kL+1$ hyperedges. Note that in G' the subgraph induced by all edges containing $\{u_1, \ldots, u_k\}$ form a subtree of T. It follows that $G^k_{n,p} \oplus T$ is at least as likely to be non-homomorphic to H as G' is, so:

$$\Pr[G^k_{n,p} \oplus M \text{ is not homomorphic to } H] \leq \Pr[G^k_{n,p} \oplus T \text{ is not homomorphic to } H] + \frac{\alpha}{2}.$$

By Lemma 5, increasing p by an additional ϵp a.s. results in the addition of a copy of T. Thus:

$$\Pr[G^k_{n,p} \oplus T \text{ is not homomorphic to } H] \leq \Pr[G^k_{n,p(1+\epsilon)} \text{ is not homomorphic to } H]$$

which yields a contradiction to Lemma 4(b).

This proves the case where $k \geq 3$, so we now turn to the case $k = 2$. If H contains no triangle, then K_3 is not homomorphic to H. Thus, K_3 forms a unicyclic unsatisfiable CSP using the H-colouring constraints and so we not have a sharp threshold. So we will focus on graphs H that contain a triangle. Our proof follows along the same lines as the case $k \geq 3$, but is complicated a bit since we can no longer assume that M is homomorphic to a single edge. We only highlight the differences.

Define M^+ to be the subgraph of $G = G^k_{n,p} \oplus M$ induced by all vertices within distance $2|V(H)| + |V(M)| + 1$ of the unique cycle of M. Again, there is some constant L such that with probability at least $1 - \frac{\alpha}{2}$: M^+ is unicyclic and has at most L vertices.

Define U to be the set of vertices of G that are of distance exactly $r = 2|V(H)| + |V(M)|$ from the unique cycle of M. Consider any vertex $u \in H$. It is an easy exercise to prove that if M^+ is unicyclic then there is a homomorphism from M^+ to H such that all vertices in U are mapped to u. We omit the proof (which only requires the fact that $r > 2|V(H)|$.)

Obtain the graph G' from G by (i) removing all of the vertices of distance less than r from the unique cycle of M, and (ii) contracting U into a single new vertex u. Suppose that h' is a homomorphism from G' to H. Then by the previous paragraph, h' can be extended to a homomorphism from G to H where each vertex $v \in V(G') - u$ is mapped to $h'(v)$, and every vertex in U is mapped to $h'(u)$. Thus, if G' is homomorphic to H then so is G.

Let T be the tree which consists of a vertex adjacent to L leaves. Since the degree of u in G' is at most L, and using the fact that all vertices of M are deleted when forming G' (here is where we require $r > |M|$), the rest now follows as in the $k \geq 3$ case. □

We close this section by discussing the possibilities of extending Theorem 3 to the case where H is disconnected. We will focus on graphs, i.e. the $k = 2$ case. We can show that if Hypothesis A does not hold for the H-homomorphism problem for every graph H, then there must be a counterexample with two components: a triangle and a graph H_1 that is triangle-free and not 3-colourable. (We omit the argument.) So the question of whether there is any graph H for which the H-homomorphism problem has a coarse threshold is equivalent to the following:

Question 7 *Is there any triangle-free graph H with $\chi(H) > 3$ such that for some values of n and some $c > c(n)$, $G_{n,p=c/n}$ is not a.s. non-H-homomorphic?*

4 The (d, k, t)-model

Proof of Theorem 2 Suppose that the (d, k, t)-model exhibits a coarse threshold. Then consider p, α, ϵ and M as guaranteed by Lemma 4. It is easy to verify that, since $t < d^{k-1}$ and M is unicyclic, M is satisfiable. Suppose that $V(M) = u_1, ..., u_r$ and let a_i be the value of u_i in some particular satisfying assignment A of M. Given a CSP F on at least r variables, we define $F \oplus A$ to be the CSP formed by choosing a random ordered r-tuple of variables $v_1, ..., v_r$ in F and for each $1 \leq i \leq r$, forcing v_i to take the value a_i by adding a one-variable constraint on v_i. Clearly the probability that $\widehat{CSP}_{n,p}(\mathcal{P}) \oplus A$ is unsatisfiable is at least as high as the probability that $\widehat{CSP}_{n,p}(\mathcal{P}) \oplus M$ is unsatisfiable.

Straightforward first moment calculations imply that there is some constant L such that, defining E_1 to be the event that every v_i has at most L neighbours, $\mathbf{Pr}(E_1) \geq 1 - \frac{\alpha}{2}$. Suppose that E_1 holds. Consider a particular v_i and expose the $\ell \leq L$ edges containing it, $e_1, ..., e_\ell$ and the corresponding constraints $C_1, ..., C_\ell$. For each C_j, let C_j' be the $(k-1)$-variable constraint obtained by restricting v_i to be a_i; i.e. a $(k-1)$-tuple of values is permitted for C_j' iff C_j permits that same $(k-1)$-tuple along with $v_i = a_i$. Since C_j has at most t restrictions, C_j' has at most t restrictions.

Let G be a random CSP formed as follows: start with a random $\widehat{CSP}_{n,p}(\mathcal{P})$ and then for each of the $\binom{d^{k-1}}{t}$ possible constraints on $k-1$ variables and with t restrictions, choose rL random ordered $(k-1)$-tuples of variables and place that constraint on them. It is easy to see that the probability that G is unsatisfiable is at least as high as the probability that $\widehat{CSP}_{n,p}(\mathcal{P}) \oplus A$ is unsatisfiable; each batch of L copies of every $(d, k-1, t)$-constraint is at least as restrictive as forcing $v_i = a_i$. Thus, adding those rL constraints boosts the probability of unsatisfiability by at least 2α. We say that a set of $\binom{d^{k-1}}{t} rL$ ordered $(k-1)$-tuples is *bad* if adding the constraints to that set results in an unsatisfiable CSP. So, consider the following random experiment: pick a random $\widehat{CSP}_{n,p}(\mathcal{P})$ and then pick $\binom{d^{k-1}}{t} rL$ ordered $(k-1)$-tuples of the variables. The probability that we pick a bad set is at least 2α. Since $\binom{d^{k-1}}{t} rL = O(1)$, a simple first moment calculation shows that a.s. the choice of $(k-1)$-tuples will be vertex disjoint. Thus, the probability of picking a bad set is at least $2\alpha - o(1)$ even if we condition on the $(k-1)$-tuples being vertex-disjoint.

Define T by: (1) taking the hypergraph consisting of a vertex v lying in $rL\binom{d^k}{t}$ edges where no other vertex lies in more than one of the edges, and (2) placing each of the $\binom{d^k}{t}$ possible (d, k, t)-constraints on rL of the edges. For each $1 \leq \delta \leq d$, let T_δ denote the collection of $(k-1)$-tuples obtained by removing v from every edge containing a constraint in which every restriction has $v = \delta$; note that $|T_\delta| = \binom{d^{k-1}}{t} rL$. By Lemma 5, there is some $\zeta = o(1/n)$ such that $\widehat{CSP}_{n,p=z}(\mathcal{P})$ a.s. contains a copy of T. Consider adding that copy of T to $\widehat{CSP}_{n,p}(\mathcal{P})$. The probability that for each $1 \leq \delta \leq d$, T_δ is a bad set is at least $(2\alpha - o(1))^d$. Note that if every T_δ is a bad set, then the resulting CSP is unsatisfiable because setting $v = \delta$ requires the set of $(k-1)$-constraints on T_δ to be enforced. Thus, the probability that $\widehat{CSP}_{n,p=p+z}(\mathcal{P})$ is unsatisfiable is at least $(2\alpha - o(1))^d$. By considering adding x copies of T, we see that the probability that $\widehat{CSP}_{n,p=p+xz}(\mathcal{P})$ is satisfiable is at most $(1 - (2\alpha - o(1))^d)^x$ which is less than α for some sufficiently large constant x. Since $z = o(1/n)$, this implies that $\mathbf{Pr}(\widehat{CSP}_{n,(1+\epsilon)p}(\mathcal{P})$ is unsatisfiable$) > 1 - \alpha$ which contradicts Lemma 4(b). $\square$

5 Binary CSP's with domain size 3

Recall that Istrate[23] (see also Creignou and Daude[11]) has proven that when the domain size $d = 2$, then Hypothesis A holds; i.e. if every unicyclic CSP is satisfiable then $CSP_{n,p}(\mathcal{P})$ has a sharp threshold. This result does not extend to $d = 3$. Consider the following example, with $d = 3, k = 2$:

Example 8 *We have two constraints. C_1 says that either both variables are equal to 1, or neither is equal to 1. C_2 says that the variables cannot both have the same value.* $\mathcal{P}(C_1) = \frac{2}{3}, \mathcal{P}(C_2) = \frac{1}{3}$.

It is an easy exercise to show that every unicyclic CSP that uses only constraints C_1, C_2 is satisfiable.

Consider any $\frac{3}{2} < c < 3$. Thus, a.s. the sub-CSP formed by the C_1 constraints has a giant component, and the sub-CSP formed by the C_2 constraints does not. We will show that $CSP_{n,p}(\mathcal{P})$ is neither a.s. satisfiable nor a.s. unsatisfiable.

To see that it is not a.s. unsatisfiable, note that the subgraph induced by the C_2 constraints is 2-colourable with probability at least some positive constant. If it is 2-colourable, then we can satisfy all the C_2 constraints by assigning every variable either 2 or 3; this will not violate any C_1 constraints.

To see that it is not a.s. satisfiable, consider the giant component, T, formed by the C_1 constraints and note that either every variable in T is assigned 1 or none of them are. A.s. at least one C_2 constraint has both variables in T, and so they cannot both be assigned 1. Thus, a.s. no variables in T can be assigned 1. This implies that if the C_2 constraints form an odd cycle using variables of T then the CSP is not satisfiable; that event occurs with probability at least some positive constant.

The main result of this section, is that when $d = 3$ and $k = 2$, if Hypothesis A fails on some model, then it has to fail for the same reason as it failed for Example 8.

Consider a CSP F where every constraint is on 2 variables. Suppose there is some constraint on variables v, u which implies that if v is assigned δ then u must be assigned γ; we say that $v : \delta$ *forces* $u : \gamma$ and denote this by $v : i \to u : j$. Moreover if there is a sequence of variables $v_1, \ldots, v_r$ and values $\delta_1, \ldots, \delta_r$ such that $v_i : \delta_i \to v_{i+1} : \delta_{i+1}$ for $i = 1, \ldots, r-1$ then we say that $v_1 : \delta_1$ forces $v_r : \delta_r$.

Theorem 9 *Consider some $\mathcal{P}$ with $d = 2, k = 3$ such that every unicyclic CSP formed from* **supp**$(\mathcal{P})$ *is satisfiable. $CSP_{n,p}(\mathcal{P})$ has a coarse threshold iff there exists a unicyclic CSP M formed from* **supp**$(\mathcal{P})$*, a value $1 \le \delta \le 3$, $p = p(n), \epsilon > 0, z > 0, b > 0$ such that:*

(a) $\mathbf{Pr}(CSP_{n,p}(\mathcal{P})$ *is satisfiable* $) > \epsilon$*;*

(b) M cannot be satisfied using only the two values other than δ;

(c) $\widehat{CSP}_{n,p}(\mathcal{P})$ a.s. has at least zn variables v such that $v : \delta \to u : \delta$ for at least bn variables u.

Thus, this explains the only ways that a model with $d = 3, k = 2$ can have a coarse threshold. We remark that this theorem does not extend to $d = 4, k = 2$ nor $d = 3, k = 3$; in both cases there are other causes for a coarse threshold.

Proof Outline We leave it to the reader to verify, using similar reasoning to that for Example 8 that conditions (a,b,c) imply a coarse threshold. The only difference here is that the set of variables u in (c) could change for different choices of v, and it is important to note that for each value δ, there is some constraint in **supp**$(\mathcal{P})$ that forbids both variables from receiving δ as otherwise $CSP_{n,p}(\mathcal{P})$ is trivially satisfiable by setting every variable equal to δ. We will outline a proof of the other direction. The proof is somewhat lengthy, so we will only highlight the key ideas. The case where some domain values are *bad* (as defined in [27]) is easily disposed of, so we assume that there are no such values.

For each variable v and each $1 \le \delta \le 3, 1 \le \gamma \le 3$ we define $F_{\delta,\gamma}(v)$ to be the set of variables u such that $v : \delta \to u : \gamma$, and we define $F_\delta(v) = \cup_{1 \le \gamma \le 3} F_{\delta,\gamma}(v)$. We can expose $F_\delta(v)$ by using a simple breadth-first search from v. This allows us to analyze the distribution of the size of $F_\delta(v)$ and $F_{\delta,\gamma}(v)$ using a standard branching-process analysis (see eg. Chapter 5 of [24]). We say that $F_{\delta,\gamma}$ *percolates* if there are constants $\zeta, \beta > 0$ such that $\mathbf{Pr}(|F_{\delta,\gamma}(v)| \ge \beta n) \ge \zeta$. It is straightforward to prove:

Claim 1: If $F_{\delta,\gamma}$ does not percolate, then for every $\xi > 0$ there is a constant L such that $\mathbf{Pr}(|F_{\delta,\gamma}(v)| \le L) > 1 - \xi$.

Claim 2: If $F_{\delta,\gamma}$ percolates, then there are constants $z, b > 0$ such that a.s. there are at least zn variables v with $|F_{\delta,\gamma}(v)| \ge \beta n$.

Claim 2 yields that $F_{\delta,\delta}$ percolates for at most one value δ or else $CSP_{n,p}(\mathcal{P})$ is a.s. unsatisfiable. Suppose that $F_{\delta,\delta}$ and $F_{\gamma,\gamma}$ both percolate. It is not hard to show there are a.s. $\Theta(n)$ variables v such that $|F_{\delta,\delta}(v)| = \Theta(n)$ and $|F_{\gamma,\gamma}(v)| = \Theta(n)$. It then follows that a.s. there is no satisfying assignment in which such a v is assigned δ or γ. Thus there are $\Theta(n)$ variables which all must take the third domain-value, μ. A.s. some constraint will appear on two of those variables that forbids (μ, μ).

It is less straightforward, but not very difficult, to prove:

Claim 3: If $F_{\delta,\gamma}$ percolates for any pair δ, γ then either (i) $F_{\delta,\delta}$ percolates or (ii) there is some μ such that $F_{\mu,\mu}$ percolates and there is a sequence of constraints in $\mathbf{supp}(\mathcal{P})$ *through which $v : \delta \to u : \mu$.*

Suppose that $\widehat{CSP}_{n,p}(\mathcal{P})$ has a coarse threshold and consider $M, \epsilon, \alpha, p = p(n)$ from Lemma 4.

Claim 4: There is a value δ such that (i) that every satisfying assignment of M must use δ on at least one variable and (ii) $F_{\delta,\gamma}$ percolates for at least one value γ.

If $F_{\delta,\delta}$ percolates then this satisfies Theorem 9. Otherwise, applying Claim 3, $F_{\mu,\mu}$ percolates and there is a sequence of constraints so that $v : \delta \to u : \mu$. Attaching that sequence to every variable of M yields a unicyclic CSP M' for which every satisfying assignment must use μ on at least one variable. Thus M', μ satisfy Theorem 9.

Suppose that Claim 4 does not hold, and consider any satisfying assignment A of M in which every value δ used is such that $F_{\delta,\delta}$ does not percolate. Suppose that M has r variables $x_1, ..., x_i$ and that A assigns a_i to x_i. Recall from Section 4 that $CSP_{n,p}(\mathcal{P}) \oplus A$ is formed by taking $CSP_{n,p}(\mathcal{P})$ and then choosing r random variables $v_1, ..., v_i$ and adding one-variable constraints that force v_i to take a_i. Clearly $\mathbf{Pr}(CSP_{n,p}(\mathcal{P}) \oplus A \text{ is unsatisfiable}) \geq \mathbf{Pr}(CSP_{n,p}(\mathcal{P}) \oplus M \text{ is unsatisfiable})$.

Expose $F = \cup_{i=1}^{r} F_{a_i}(v_i)$, and U, the set of variables outside of F that lie in a constraint with a variable in F. Since none of the F_{a_i} percolate, Claim 2 allows us to show that there is some L such that with probability at least $1 - \alpha/2$, $|U| < L$. Since adding M to $\widehat{CSP}_{n,p}(\mathcal{P})$ increases the probability of unsatisfiability by at least 2α, it must be that the probability that $\widehat{CSP}_{n,p}(\mathcal{P})$ is satisfiable, $|U| \leq L$ and $\widehat{CSP}_{n,p}(\mathcal{P}) \oplus A$ is unsatisfiable is at least $3\alpha/2$.

Suppose that $\widehat{CSP}_{n,p}(\mathcal{P})$ is satisfiable and $|U| \leq L$. Consider some $u \in U$ sharing a constraint with $w \in F$ where A forces w to take the value μ. Let $\Omega = \Omega(u)$ be the set of values which can be assigned to u which, in conjunction with assigning μ to w do not violate their constraint. We know that $|\Omega| \neq 0$ since otherwise μ is a *bad value*. We know that $|\Omega| \neq 1$ since otherwise $u \in F$. So $|\Omega(u)| \geq 2$ for each $u \in U$. Suppose that $u_1, ..., u_\ell$ are the variables in U with $|\Omega| = 2$, and let δ_i be the value not in $\Omega(u_i)$. Consider taking a random CSP formed as follows: first take a $\widehat{CSP}_{n,p}(\mathcal{P})$ and then choose ℓ random variables $u_1, ..., u_\ell$ and force u_i to not take value δ_i using a one-variable constraint. We have proved is that the one-variable constraints boost the probability of unsatisfiability by $3\alpha/2$.

At this point, we can complete the proof using the argument from the Achlioptas-Friedgut proof[3] that d-colourability has a sharp threshold for $d \geq 3$. In that paper, the starting point was to note that since every constraint is a colourability constraint, fixing an assignment on M at worst forbids one colour from each neighbour of M. This put them in the same position that we are in now, and so the rest of our proof is the same as theirs. □

We close this section by noting why this proof can't be extended to general d. The problem is that possibly some of the variables in U would have their domain sizes reduced by two instead of one. The argument in [3] cannot handle that possibility.

6 Future Directions

There is clearly much work still to be done along these lines of research. The big problem still remains - determine precisely which models from [27] have a sharp threshold. Of course, Section 2 indicates that this may be overly ambitious. But lowering our sights only slightly, we can try to determine all possible *causes* for coarse thresholds, i.e. continue the course started in Section 5. An important subgoal would be to do this for binary CSP's, i.e. the case where $k = 2$. Another reasonable goal to pursue would be to cover the $d = 3$ case.

As far as more specific classes of models go, one should try to extend the work in Section 3 and examine whether Hypothesis A holds for H-homomorphism problems when H is a *directed* hypergraph. Such homomorphism problems are equivalent to CSP's in which every constraint is identical under some permutation of the variables. And of course, it would be good to determine whether the "connected" condition can be removed from Theorem 3.

References

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

[2] D. Achlioptas, A. Chtcherba, G. Istrate and C. Moore. *The phase transition in NAESAT and 1-in-k SAT.* Proceedings of SODA 01, 721 - 722.

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

[4] 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.

[5] D. Achlioptas and C. Moore. *The asymptotic order of the k-SAT threshold* Proceedings of FOCS 2002.

[6] D. Achlioptas and C. Moore. *Almost all graphs of degree 4 are 3-colorable.* Proceedings of STOC 2002.

[7] V. Chvátal and B. Reed. *Mick gets some (the odds are on his side).* Proceedings of FOCS 1992, 620 - 627.

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

[9] H. Connamacher and M. Molloy. *The exact satisfiability threshold for a potentially intractible random constraint satisfaction problem.* Proceedings of FOCS 2004.

[10] N. Creignou and H. Daude. *Generalized satisfiability problems: minimal elements and phase transitions.* Theoretical Computer Science **302** (2003), 417 - 430. Preliminary version in proceedings of SAT 2002.

[11] N. Creignou and H. Daud *Combinatorial sharpness criterion and phase transition classification for random CSPs.* Information and Computation **190** (2004), 220 - 238.

[12] O. Dubois and J. Mandler. *On the non-3-colourability of random graphs.* submitted.

[13] O. Dubois and J. Mandler. *The 3-XORSAT Threshold.* Proceedings of FOCS 2002, 769 - 778.

[14] P. Erdos and A. Renyi. *On random graphs I.* Publ. Math. Debrecen **6** (1959), 290 - 297.

[15] A. Flaxman. *A sharp threshold for a random constraint satisfaction problem.* preprint (2003).

[16] N. Fountoulakis and C. McDiarmid. *Upper bounds on the non-3-colourability threshold of random graphs.* Disc. Math. and Th. Comp. Sci. **5** (2002), 205 - 226.

[17] 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.

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

[19] E. Friedgut. Personal communication.

[20] L. Goldberg, S. Kelk and M. Patterson. *The complexity of choosing an H-colouring (nearly) uniformly at random.* Proceedings of STOC 2002.

[21] P. Hell and J. Nesetril. *Graphs and Homomorphisms.* Oxford University Press (2004).

[22] P. Hell and J. Nesetril. *On the complexity of H-colouring.* J. Comb. Th. (B) **48** (1990), 92 - 110.

[23] G. Istrate. *Phase transitions and all that.* Manuscript (2002).

[24] S. Janson, T. Łuczak and A. Ruciński. Random Graphs. Wiley, New York (2000).

[25] S. Janson, Y. Stamatiou and M. Vanvakari. *Bounding the unsatisfiability threshold of random 3-SAT.* Random Structures and Algorithms **17** (2000), 103 - 116.

[26] D. Mitchell. *Resolution complexity of random constraints.* Proceedings of Principles and Practices of Constraint Programming - CP 2002.

[27] M. Molloy. *Models for random constraint satisfaction problems.* Proceedings of STOC 2002, 209 - 217.

[28] M. Molloy. *Thresholds for colourability and satisfiability for random graphs and boolean formulae.* In "Surveys in Combinatorics", J. Hirschfield, ed., Cambridge University Press, 2001, 165 - 197.

[29] M. Molloy. *When does the giant component bring unsatisfiability?* Submitted.

[30] M. Molloy and M. Salavatipour. *The resolution complexity of random constraint satisfaction problems.* Proceedings of FOCS 2003. (Journal version submitted.)

[31] M. Talagrand. *On Russo's approximate zero-one law.* Ann. of Prob. **22** (1994), 1576 - 1587.

[32] J. van den Berg and H. Kesten. *Inequalities with application to percolation and reliability.* J. App. Prob. **22** (1985), 556 - 569.

[33] K. Xu and W. Li. *Exact Phase Transitions in Random Constraint Satisfaction Problems.* J. Artificial Intelligence Research **12**, 93 - 103 (2000).

7 Appendix - Proof of Lemma 4

This notion of sharp thresholds was introduced by Friedgut in [17], where he provided a technique to prove the existence of sharp thresholds. Friedgut's main theorems consider random graphs. The aim of this appendix is to modify these theorems so that they can be applicable to random CSP's. Instead of being restricted to the unsatisfiability property, we consider, as Friedgut did, every monotone property where a property A is called *monotone* if it is preserved under constraint addition. A property A on CSP's is called *monotone symmetric* if it is monotone and invariant under CSP automorphisms. For a property A, A_i denotes the restriction of A on CSP's with exactly n variables.

Lemma 4 is a corollary of the following theorem which can be applied to a wider range of problems.

Theorem 10 *Let $A = \{A_i\}$ be a series of monotone symmetric properties in $\widehat{CSP}_{n,p}(\mathcal{P})$ with a coarse threshold. There exist, $p = p(n)$, $\tau, \alpha, \epsilon > 0$, a CSP M whose constraints are chosen from* $\mathbf{supp}(\mathcal{P})$ *such that for an infinite number of n:*

(a) $\alpha < \Pr[\widehat{CSP}_{n,p}(\mathcal{P}) \in A] < 1 - 3\alpha$.

(b) $\tau < \Pr[M \subseteq \widehat{CSP}_{n,p}(\mathcal{P})] < 1 - \tau$.

(c) $\Pr[\widehat{CSP}_{n,p}(\mathcal{P}) \oplus M \in A] > 1 - \alpha$.

(d) $\Pr[\widehat{CSP}_{n,p(1+\epsilon)}(\mathcal{P}) \in A] < 1 - 2\alpha$.

Given two properties A and A', their symmetric difference, $A\Delta A'$ is the property of satisfying one but not both of them. Given a set of CSP's $\mathcal{H} = \{H_1, \ldots, H_k\}$, we define $S(\mathcal{H})$ to be the property of having a subproblem which is isomorphic to one of the members of $\mathcal{H}$. Let μ_p be the probability that $\widehat{CSP}_{n,p}(\mathcal{P}) \in A$, and let $I_{p^*} = \frac{d_\mu}{d_p}|_{p=p^*}$.

Suppose that a series of monotone properties A has a coarse threshold. Then since the conditions of the sharp threshold are not satisfied, there exists $\epsilon_{\rm thr} < \frac{1}{2}$ and $C > 0$ such that if p_0 is defined as $\mu_{p_0,n}(A_n) = \epsilon_{\rm thr}$ and p_1 is defined as $\mu_{p_1,n}(A_n) = 1 - \epsilon_{\rm thr}$, then $\frac{p_1}{p_0} > 1 + C$ for infinitely many n. Moreover there is a function $p = p(n)$ such that for infinitely many n, $p.I < c$ for a constant c.

A *balanced hypergraph* is a hypergraph with average degree no smaller than that of any of its subgraphs. A *strictly balanced hypergraph* is one where the average degree is strictly larger than that of any proper subgraph. A CSP is call *balanced* (or *strictly balanced*) if its constraint hypergraph is balanced (or strictly balanced).

The following two theorems are analogues of the main theorems of Friedgut's paper [17].

Theorem 11 *There exists a function $k(\epsilon, c)$, such that for all $c > 0$, any n and any monotone symmetric family of CSP's Q on n variables, such that $p.I \leq c$, for every $\epsilon > 0$ there exists a set $\mathcal{H} = \{H_1, \ldots, H_r\}$ of CSP's such that $|E(H_i)| \leq k(\epsilon, c)$ for $1 \leq i \leq r$ and $\mu_p(Q\Delta S(\mathcal{H})) \leq \epsilon$. Furthermore the constraint hypergraph of H_i are all balanced.*

Theorem 12 *Let $0 < \epsilon_{\rm thr} < \frac{1}{2}$. There exist functions $B(\alpha, c), b_1(\alpha, c), b_2(\alpha, c)$ such that for all $c > 0$, any n and any monotone symmetric family of hypergraphs Q on n variables such that $p.I \leq c$ and $\epsilon_{\rm thr} < \mu_p(Q) < 1 - \epsilon_{\rm thr}$, for every $\alpha > 0$ there exists a hypergraph M with the following properties.*

- *The constraint hypergraph of M is balanced.*
- $b_1 < \mathbf{Ex}[M] < b_2$.
- $|E(M)| \leq B$.
-
$$\Pr[\widehat{CSP}_{n,p}(\mathcal{P}) \oplus M \in Q] \geq 1 - \alpha.$$

Although Theorem 11 provides a better intuition of the coarse thresholds, Theorem 12 is more suitable for proving that certain properties have a coarse threshold. The following lemma shows that how Theorem 12 implies Theorem 10.

Lemma 13 *Let $A = \{A_i\}$ be a series of monotone properties of $\widehat{CSP}_{n,p}(\mathcal{P})$ with a coarse threshold, i.e. $\frac{p_1}{p_0} > 1 + C$ for infinite number of n. There exists a function $\epsilon(\alpha)$ which satisfies the property that for every $\alpha > 0$ and infinitely many n, there exists some $p_k(\alpha, n)$ such that defining $p'_k = p_k(1+\epsilon)$, we have*

- $p_k I_{p_k} < \frac{(2+C)}{C} = C'$.
- *We have* $\epsilon_{\text{thr}} < \mu_{p_k}(A) \leq \mu_{p'_k}(A) < 1 - \epsilon_{\text{thr}}$.
- $\mu_{p'_k}(A) < \mu_{p_k}(A) + \alpha$.

Proof. Let p'_1 be defined as $p'_1 = (1+C)p_0$. Obviously $p'_1 \leq p_1$ for infinitely many n. Let

$$R = \{p \in [p_0, p'_1] : p.I_p \leq C'\},$$

and $R' = [p_0, p'_1] - R$. For fixed n, $\mu_{n,p}(A_n)$ is a polynomial of degree $N = k!\binom{n}{k}|\mathbf{supp}(\mathcal{P})|$ in p. So R is a finite union of intervals.

$$1 > \mu_{p'_1} - \mu_{p_0} > C' \int_{R'} \frac{1}{x} dx > \frac{C'}{p'_1} \int_{R'} dx > \frac{C'(p'_1 - p_0 - |R|)}{p'_1} \geq \frac{2+C}{1+C} - \frac{|R|}{(C+1)p_0},$$

So we have $|R| > p_0$. Let $\epsilon = \frac{\alpha}{4}$. There exists a sequence $\hat{p}_0 < \ldots < \hat{p}_{\lceil \frac{2}{\alpha} \rceil}$ where $\hat{p}_i \in R$ and no pair of these values are in a distance of less than ϵp_c. Obviously there exists some $1 \leq i < \lceil \frac{2}{\alpha} \rceil$ such that $\mu_{\hat{p}_{i+1}}(A) < \mu_{\hat{p}_i}(A) + \alpha$.

Since all of these points are in R, we have $\hat{p}_i.I_{\hat{p}_i} < C'$. ■

7.1 Proofs of Main Theorems

The proof of Theorems 11 and 12 are parallel to the proofs of Friedgut's original theorems [17]. In fact we verify that every step of the proofs in [17] holds for CSP's.

In what follows, by a CSP we mean a CSP on n variables whose constraints are in $\mathbf{supp}(\mathcal{P})$. For CSP's H and S we will denote the fact that they are isomorphic by $H \sim S$. Let E be the set of all $|\mathbf{supp}(\mathcal{P})|\binom{n}{k}k!$ possible appearance of the constraints in $\mathbf{supp}(\mathcal{P})$ placed on k variables which are chosen from the set of n variables. The complete CSP on n variables is the CSP which contains all constraints in E. For every $e \in E$, let $p_e = \mathcal{P}(e).p/k!$. So e is present in $\widehat{CSP}_{n,p}(\mathcal{P})$ with probability

p_e and independent of all other elements of E. Define $q_e = 1 - p_e$. For a CSP H, let $\Theta(H)$, the orbit of H be the set of all subsets of the complete CSP on n vertices which are isomorphic to H. Note that $\mathbf{Ex}[H] = |\Theta(H)| \prod_{e \in H} p_e \leq |\Theta(H)| p^{|H|}$. Let $D(H) = n^{v(H)} p^{|H|}$, where $v(H)$ denotes the number of vertices of H. Note that for H of constant size $D(H) \geq \mathbf{Ex}[H] \geq cD(H)$.

For any monotone set $A \subseteq \{0,1\}^n$, define

$$||A|| = \max\{\sum \epsilon_i : \epsilon \text{ is a minimal element in } A\}.$$

7.1.1 Fourier analysis

We define an orthogonal basis, with respect to μ, for the space of real functions on $\widehat{CSP}_{n,p}(\mathcal{P})$.

For every $e \in E$, let U_e be defined as follows:

$$U_e(H) = \begin{cases} -\sqrt{\frac{p_e}{q_e}} & \text{if } e \in H \\ \sqrt{\frac{q_e}{p_e}} & \text{if } e \notin H \end{cases}$$

For every CSP R, define

$$U_R = \prod_{e \in R} U_e.$$

Define $U_\emptyset$ to be identically equal to 1. These functions indeed are orthonormal with respect to the inner product defined by μ. For every real function f on $\widehat{CSP}_{n,p}(\mathcal{P})$, and every CSP H, define $\hat{f}(H) = \langle f, U_H \rangle$. This gives us the Fourier expansion of f:

$$f = \sum_H \hat{f}(H) U_H.$$

For any value of p, the usual norm of f is defined as $||f||_2^2 = \langle f, f \rangle$. We have

$$||f||_2^2 = \sum \hat{f}^2.$$

For a given $e \in E$, define $A_e(f)$, the influence of e on f, to be the set of CSP's H such that $f(H) \neq f(H \setminus e)$. Put $I = \frac{1}{p} \sum_{e \in E} \mu(A_e)$. Let $f = \chi(A)$ be the characteristic function of the monotone property A.

The following lemma can be easily proved by induction.

Lemma 14 $d_\mu(A)/d_p = I$.

The next lemma shows the relation between Fourier coefficients and $\mu(A_e)$.

Lemma 15 $q_e \mu(A_e) = \sum_{H : e \in H} \hat{f}^2(H)$. *Consequently* $p.I \geq \sum_H |H| \hat{f}^2(H)$.

Proof. The proof is similar to [31]. Let $H \ominus e$ be the CSP obtained from H by deleting e if e is a constraint of H, and adding it if it is not. For every constraint e, define

$$f_e(H) = \begin{cases} q_e(f(H) - f(H \ominus e)) & \text{if } f(H) = 1, \\ p_e(f(H) - f(H \ominus e)) & \text{if } f(H) = 0. \end{cases} \tag{1}$$

It is not hard to verify that

$$\hat{f}_e(H) = \begin{cases} \hat{f}(H) & \text{if } e \in H, \\ 0 & \text{if } e \notin H. \end{cases} \tag{2}$$

So we have

$$f_e = \sum_{H: e \in H} \hat{f}(H) U_H.$$

This implies that

$$||f_e||_2^2 = \sum_{H: e \in H} \hat{f}^2(H). \tag{3}$$

The result follows from that fact that $||f_e||_2^2 = q_e \mu(A_e)$. ■

7.1.2 Some Lemmas on random CSP's

For every CSP S we define:

$$V = V_S = \sum_{H \in \Theta(S)} U_H.$$

For a CSP R, let X_R be the random variable counting the number of copies of R in $\widehat{CSP}_{n,p}(\mathcal{P})$. Let $X_\emptyset$ be defined to be identically 1.

Lemma 16 *For any fixed CSP, S.*

$$V_S = \left(\prod_{e \in S} \sqrt{\frac{1}{p_e q_e}} \right) \mathbf{Ex}[S] \left(\sum_{R \subseteq S} (-1)^{|R|} \frac{X_R}{\mathbf{Ex}[R]} \right)$$

Proof. Let H be a fixed copy of S, and for every $H' \subseteq H$ let $Y_{H'}$ be an indicator random variable taking the value 1 if and only if all constraints of H' appear in $\widehat{CSP}_{n,p}(\mathcal{P})$. $U_H(\widehat{CSP}_{n,p}(\mathcal{P}))$ is determined by the maximal H' such that $Y_{H'} = 1$.

$$U_H(\widehat{CSP}_{n,p}(\mathcal{P})) = \left(\prod_{e \in H} -\sqrt{q_e/p_e} \right) \left(\prod_{e \in H - H'} -p_e/q_e \right).$$

This can be expressed as follows:

$$\begin{aligned} & U_H(\widehat{CSP}_{n,p}(\mathcal{P})) = \\ & \left(\prod_{e \in H} -\sqrt{\frac{q_e}{p_e}} \right) \left[\sum_{H_1 \subseteq H} \left(\prod_{e \in H - H_1} -\frac{p_e}{q_e} \right) \sum_{H_1 \subseteq H_2 \subseteq H} (-1)^{(|H_2| - |H_1|)} Y_{H_2} \right] = \\ & \left(\prod_{e \in H} \sqrt{\frac{q_e}{p_e}} \right) \left[\sum_{H_2 \subseteq H} (-1)^{|H_2|} Y_{H_2} \sum_{H_1 \subseteq H_2} \left(\prod_{e \in H - H_1} \frac{p_e}{q_e} \right) \right] \end{aligned} \tag{4}$$

In calculating V_S we sum (4) on all $H \in \Theta(S)$. Using $X_H = \sum Y_H$, we get:

$$V_S = \left(\prod_{e \in H} \sqrt{\frac{q_e}{p_e}}\right) \left[\sum_{H_2 \subseteq H} (-1)^{|H_2|} \frac{|\Theta(H)|}{|\Theta(H_2)|} X_{H_2} \sum_{H_1 \subseteq H_2} \left(\prod_{e \in H - H_1} \frac{p_e}{q_e}\right)\right]$$

Using

$$\prod_{e \in H - H_1} \frac{p_e}{q_e} = \left(\prod_{e \in H - H_2} \frac{p_e}{q_e}\right) \left(\prod_{e \in H_2 - H_1} \frac{p_e}{q_e}\right),$$

and the easy equation

$$\sum_{H_1 \subseteq H_2} \left(\prod_{e \in H_2 - H_1} \frac{p_e}{q_e}\right) = \prod_{e \in H_2} \frac{1}{q_e},$$

we get

$$V_S = \left(\prod_{e \in S} \sqrt{\frac{q_e}{p_e}}\right) \left[\sum_{R \subseteq S} (-1)^{|R|} \frac{|\Theta(S)|}{|\Theta(R)|} X_R \left(\prod_{e \in S - R} \frac{p_e}{q_e}\right) \prod_{e \in R} \frac{1}{q_e}\right].$$

Thus

$$V_S = \left(\prod_{e \in S} \sqrt{\frac{1}{p_e q_e}}\right) \left[\sum_{R \subseteq S} (-1)^{|R|} X_R \frac{|\Theta(S)| \prod_{e \in S} p_e}{|\Theta(R)| \prod_{e \in R} p_e}\right].$$

■

Let e be a constraint in E. For any CSP H, we define

$$V_{H,e} = \sum_{R \in \Theta(H), e \in R} U_R.$$

Note that

$$V_S = \frac{1}{|S|} \sum_e V_{S,e}.$$

The expected value of $V_{S,e}$ is its inner product to the identity $U_\emptyset$ which is 0 as they are orthogonal.

Let $\tilde{H}$ be a non-empty subproblem of H such that $D(\tilde{H})$ is minimal, i.e.

$$\forall R \subseteq H, D(\tilde{H}) \leq D(R).$$

Note that $\tilde{H}$ must be an induced subproblem.

Lemma 17 $\Pr\left[|V_{H,e}| \geq \lambda \sqrt{\Theta(H)} / \left(D(\tilde{H})^{1/4}\right)\right] \leq c\lambda^{-4}/n^k$.

Proof. The lemma will follow from a bound on the 4th moment of $V_{H,e}$.

$$\mathbf{Ex}[V_{H,e}^4] = \mathbf{Ex}\left[\left(\sum_{R \in \Theta(H), e \in R} U_R\right)^4\right] \leq c\mathbf{Ex}\left[\sum_S |\Theta(S)| n^{-k} \sum_L \prod_{R \in L} U_R\right], \tag{5}$$

where on the right hand side the first sum is over all isomorphism types of CSP's S that S is a union of 4 copies of H having the constraint e in their intersection. The second sum is on L's that are quadruples of copies of H with a constraint e common to all 4 of them. n^-2 is appeared on the righthand side because S must contain the fixed constraint e.

Let us compute the contribution of a given S to the sum. Assume S is the union of 4 copies of H. Let us compute the contribution of a given S to the sum. Assume S is the union of 4 copies of H. Let c_e be the number of times a constraint e of S is covered by these copies.

$$\mathbf{Ex}[\prod U_R] = \prod \mathbf{Ex}[U_e^{c_e}] = \prod \left(p_e \left(-\sqrt{q_e/p_e} \right)^{c_e} + q_e \left(\sqrt{p_e/q_e} \right)^{c_e} \right).$$

We have used the fact that for different constraints $e \neq e'$ the values of U_e and U'_e are independent random variables. Now if $c_e = 1$ for some constraints then $\mathbf{Ex}[\prod U_R] = 0$. Otherwise the dominant summand corresponding to e is $p_e(-\sqrt{q_e/p_e})^{c_e}$ and since S has constant size:

$$\mathbf{Ex}[\prod U_R] \sim p^{|S|}/\sqrt{p}^{(4|H|)}. \tag{6}$$

Note that the number of summands in the double sum $\sum_S \sum_L$ is a constant depending on H. Substituting $\mathbf{Ex}[S] \geq c'|\Theta(S)|p^{|S|}$, for some constant c, and (5) and (6) we obtain

$$\mathbf{Ex}[V_{H,e}^4] = cp^{-2|H|} \max_S \mathbf{Ex}[S]/n^2 \tag{7}$$

where the maximum is only taken over CSP's S that can be covered by 4 copies of H with a constraint in common, each constraint being covered at least twice. Recall that $\mathbf{Ex}[G] \leq D(G)$. Therefore we have

$$\mathbf{Ex}[V_{H,e}^4] = cp^{-2|H|} \max_S D(S)/n^2 \tag{8}$$

claim: The S for which $D(S)$ is maximized among the CSP's in question consists of the union of 2 copies of H, whose intersection is exactly $\tilde{H}$. For this S, we have $D(S) = n^{2v(H)-v(\tilde{H})}n^p 2v(H) - v(\tilde{H})$. Then (8) will give

$$\mathbf{Ex}[V_{H,e}^4] = cn^{2v(H)-k}/D(\tilde{H}).$$

Recall that $|\Theta(H)| \approx cn^{v(H)}$, and the desired result follows from Markov's inequality. ■

Corollary 18 *Let*

$$\chi = \chi \left\{ |V_{H,e}| > \sqrt{\Theta(H)}/ \left(D(\tilde{H})^{1/4} \right) \right\}.$$

Then

$$\int |V_{H,e}|.\chi \leq c(\sqrt{\Theta(H)}/ \left(D(\tilde{H})^{1/4} \right) /n^k.$$

Lemma 19 *Let R be a fixed CSP, and let X_R be a random variable equal to the number of copies (not necessarily induced copies) of R that appear in a random CSP $\widehat{CSP}_{n,p}$. Assume $p = o(1)$. Then*

$$\mathbf{Var}[X_R] \asymp \mathbf{Ex}[R]^2 \sum_{H \subseteq R} 1/\mathbf{Ex}[H]$$

where the sum is over all non-empty subproblems of R.

Proof. $\mathbf{Var}[X_R]$ is given by the formula

$$\sum \mathbf{Ex}[XY] - \mathbf{Ex}[X]\mathbf{Ex}[Y]$$

where the sum is over all pairs (X, Y) that are indicator random variables indicating whether a specific copy of R appeared in the random CSP. If X and Y are independent the corresponding summand is 0, otherwise $\mathbf{Ex}[X]\mathbf{Ex}[Y] = o(\mathbf{Ex}[XY])$. Partitioning the sum $\sum \mathbf{Ex}[XY]$ according to the isomorphism type of the intersection of the two copies of R indicated by X and Y gives the result. ■

7.1.3 Main Lemma

Let H be a CSP. Choosing certain bounds L, c_1, c_2 call H modest if

- $|H| \le L$.
- $c_1 \le \mathbf{Ex}[H] \le c_2$.
- H is balanced.

If there exists a CSP H of bounded size such that $\mathbf{Ex}[H]$ is a constant, then there exists a constant c such that for any other CSP G of bounded size and any fixed m, if n is large enough, then

$$\mathbf{Ex}[G] < \log(n)^m \Rightarrow \mathbf{Ex}[G] \le c.$$

This is true since the fact that $\mathbf{Ex}[H]$ is a constant implies that $p = n^{-v(H)/|H|}$. We will use this in the proof of the following lemma.

Lemma 20 *Let A be a monotone symmetric family of CSP's on n vertices, and let f be its characteristic function. Assume that for $\widehat{CSP}_{n,p}(\mathcal{P})$, we have $p.I \le c$. Then for every $\epsilon > 0$, there exist constants L, c_1, c_2 such that for sufficiently large n.*

$$\sum_{S \text{ is not modest}} \hat{f}^2(S) \le \epsilon.$$

The rest of this section is devoted to the proof of this lemma. It follows from Lemma 15 that

$$\sum_{H:|H|>L} \hat{f}(H) \le pI/L.$$

So for every ϵ, there exists sufficiently large L such that$\sum_{H:|H|>L} \hat{f}(H) \le \epsilon$. So the most of the weight of $\hat{f}$ is concentrated on CSP's H with no more than a certain constant number of constraints. We now wish to further characterize the CSP's H such that $\sum_{S \sim H} \hat{f}^2(S)$ is significant. A simple calculation shows that for any Boolean function f and any CSP S:

$$\hat{f}^2(S) \le (4p)^{|S|}. \tag{9}$$

Summing this on the orbit of S we get

$$\sum_{H \sim S} \hat{f}^2(H) \le \mathbf{Ex}[H]. \tag{10}$$

So if H has small expectation, then its orbit does not contribute much to the weight of $\hat{f}^2$. The following lemma shows this is true even if H has a subproblem of small expectation.

Lemma 21 *Let H be a CSP. Then for every subproblem H'*

$$\sum_{S \in \Theta(H)} \hat{f}^2(S) \le c. \max_{R \subset H'} \frac{\mathbf{Ex}[H']}{\mathbf{Ex}[R]}$$

where the expectation of the empty CSP is taken to be 1

Proof. Let R be a specific copy of H'. Consider the probability space $\{0,1\}^{E\setminus R}$, where $E\setminus R$ is the set of constraints of the complete CSP not in R. We view this as the space of random CSP's on n vertices where one copy of H' is fixed (chosen with probability 1), and all other constraints e are chosen with usual probability p_e. We define the set of functions $\{U_S\}$ as before and have a Fourier expansion of any real function on this space.

Define a function g on this space by $g(G) = f(G\cup R)$. Note that g is symmetric in the sense that it is invariant under automorphisms of the complete CSP that keep R fixed. We want to show that

$$\hat{g}(G) = \sum_{R'\subseteq R} \hat{f}(R'\cup G)U_{R'}(R). \tag{11}$$

Assume first that $R=e$, R is a single constraint. In this case

$$\hat{G} = \sum_{e\notin T} f(T\cup e)U_G(T)\left(\prod_{e'\in T} p_{e'}\right)\left(\prod_{e'\notin T\setminus e} q_{e'}\right).$$

The right hand side of (11) is

$$\hat{f}(G) - \sqrt{q_e/p_e}\hat{f}(G\cup e) = \sum_M f(M)\left(U_G(M) - \sqrt{q_e/p_e}U_{G\cup e}(M)\right)\left(\prod_{e'\in M} p_{e'}\right)\left(\prod_{e'\notin M} q_{e'}\right).$$

Now for M such that $e\notin M$, the coresspondin summand is zero. Setting $M = T\cup e$, and using the fact that $U_{G\cup e} = -(\sqrt{q/p})U_G(M\setminus e)$ gives the desired result.

For $|R|>1$, pick $e\in R$ and define $g(G) = h(G\cup e)$, where h is a function on $\{0,1\}^{E\setminus(R\setminus e)}$. We already know that

$$\hat{g}(G) = \hat{h}(G) - (\sqrt{q_e/p_e})\hat{h}(G\cup e)$$

and the result follows by using the induction hypothesis on h.

For any G in this new space define $\Theta(G)$ to be the new orbit of G, under the action of automorphisms of the complete CSP that keep R fixed. Since g is Boolean:

$$|\hat{g}(G)| \le \frac{1}{\sqrt{|\tilde{\Theta}(G)|}}.$$

Using (11) we have for any CSP G such that its constraint set is disjoint from R:

$$\left|\hat{f}(R\cup G)/\sqrt{\prod_{e\in R} p_e}\right| \le c\left(1/\sqrt{|\tilde{\Theta}(G)|} + \sum_{R'\subset R}\left|\hat{f}(R'\cup G)\sqrt{\prod_{e\in R'} p_e}\right|\right) \le \tag{12}$$

$$c\left(1/\sqrt{|\tilde{\Theta}(G)|} + \sum_{R'\subset R}\left(\sqrt{\prod_{e\in R'} p_e}\sqrt{\Theta(R'\cup G)}\right)^{-1}\right). \tag{13}$$

Note that for $R'\subseteq R$,

$$|\Theta(R'\cup G)| \ge c|\Theta(R')|.|\tilde{\Theta}(G)|. \tag{14}$$

The value of c in the preceding inequality may depend on the CSP's involved, however in our vase we will be using this inequality for a finite number of CSP's, hence it holds with a fixed c. Observe that when $v(R') \cap v(G) = v(R) \cap v(G)$, the two sides of (14) are comparable, i.e.

$$|\Theta(R' \cup G)| \leq c'|\Theta(R')|.|\tilde{\Theta}(G)|.$$

Let G be such that $G \cup R = H$. Multiplying both sides of (13) by $\sqrt{\Theta(H) \prod_{e' \in R} p_{e'}}$ and using (14) gives

$$\sqrt{\sum_{S \in \Theta(H)} \hat{f}^2(S)} \leq c \left(\sqrt{\mathbf{Ex}[R] + \sum_{R' \subset R} \sqrt{\mathbf{Ex}[R]\mathbf{Ex}[R']}} \right)$$

which gives the desired result. ■

Assume that H is such that $W = \sum_{R \in \Theta(H)} \hat{f}^2(R)$ is bounded away from zero by some constant. We already know that H is of bounded size, and with expectation bounded from 0. We will further characterize H, by examining a few cases, and thus prove the following claim:

Lemma 22 *Under the above conditions H must be balanced, and of bounded expectation.*

Proof. Remember the definition of functions f_e in (1). Let S be the set of all G such that $f(G) = 0$ and $f(G \cup e) = 1$. We have

$$||f_e||_2^2 = \sum_{G \in S} p_e^2 \mu(G) + \sum_{G \in S} q_e^2 \mu(G \cup e) = \sum_{G \in S} p_e \mu(G).$$

On the other hand

$$||f_e||^1 = \sum_{G \in S} p_e \mu(G) + \sum_{G \in S} q_e \mu(G \cup e) = \sum_{G \in S} 2p_e \mu(G).$$

So we have

$$||f_e||_2^2 = \frac{1}{2}||f_e||_1.$$

and by Lemma (15), we have

$$p.I \geq c \sum_e \int |f_e|. \tag{15}$$

Recall that $\tilde{H}$ was defined as a non-empty subproblem of H on which the function D was minimal, and also $W = \sum_{R \in \Theta(H)} \hat{f}^2(R)$ and $\Theta = \Theta(H)$.

For any $R \in \Theta$, $\hat{f}(R) = \sqrt{\frac{W}{|\Theta|}}$. We will now proceed to analyze $\hat{f}$ by looking at the following expansion of f:

$$f = \sum_H \hat{f}(H) V_H = \sum_H \frac{1}{|H|} \sum_e \hat{f}(H) V_{e,H}.$$

Using the orthogonality of the function U_H we have that for any two functions f, g

$$\int f.g = \sum \hat{f}.\hat{g},$$

and in particular using symmetry and the definition of f_e

$$\int f_e.\hat{f}(H)V_{e,H} = \int (\sum_{H:e \in H} \hat{f}(H))\hat{f}(H)V_{e,H} = \sum_{R \in \Theta, e \in R} \hat{f}^2(H);$$

so,

$$W = (1/|H|) \sum_e \int f_e \hat{f}(H) V_{e,H} = (1/|H|) \sqrt{\frac{W}{|\Theta|}} \sum_e \int f_e V_{e,H}. \tag{16}$$

Let

$$\chi = \chi\{V_{e,H} > \sqrt{\Theta(H)}/D(\tilde{H})^{1/4}\}.$$

Then (16) is equal to

$$(1/|H|) \sqrt{\frac{W}{|\Theta|}} \sum_e \int f_e V_{e,H} \chi + \int f_e V_{e,H} (1-\chi).$$

Using Corollary 18, the fact that $|f_e| \leq 1$, and (15) we get that this is bounded by

$$c_1 (1/|H|) \sqrt{\frac{W}{|\Theta|}} \left(\sum_e (c_2 + pI) \sqrt{\Theta(H)}/D(\tilde{H})^{1/4} \right) = cpI\sqrt{W}/D(\tilde{H})^{1/4}.$$

So

$$D(\tilde{H}) \leq c(pI)^4/W^2. \tag{17}$$

We now assume that for a given H, $W = \sum \hat{f}^2 > \tau > 0$, for some constant $\tau > 0$ and wish to prove that H is balanced and $\mathbf{Ex}[H]$ is bounded. We will divide this into 3 cases:

Case 1: $\tilde{H}$ consists of a single constraint. In this case (17) implies that $n^k p^k \leq c'$. For such values of p, the constraint hypergraph of $\widehat{CSP}_{n,p}(\mathcal{P})$ is almost surely a disjoint union of constraints, Lemma 20 and Theorems 11 and 12 hold. The CSP property in question maybe approximated by the property of having at least k_i constraints of the ith element of $\mathbf{supp}(\mathcal{P}_\rangle)$ for some $k_i \approx (pI)^2$, and the only CSP's H of bounded size for which $\mathbf{Ex}[H]$ is not $o(1)$ are those whose constraint hypergraphs are matchings, so that $\mathbf{Ex}[H]$ is also bounded from above.

Case 2: $\tilde{H} = H$.

In this case (17) shows that $\mathbf{Ex}[H]$ is bounded by a constant, and Lemma 21 shows H must be balanced since it has no subproblems of small expectation.

Case 3: $\tilde{H} \neq H$ but $\tilde{H}$ is not a single constraint.

As before $\tilde{H}$ must be balanced and of expectation that is approximately constant. We will assume that $\tilde{H}$ is strictly balanced. If it is only balanced but not strictly balanced, replace $\tilde{H}$ by a strictly balanced subproblem of it. (Since this subproblem has the same average degree its expectation is a power of the expectation of $\tilde{H}$, and hence also bounded from above and below.)

Choose a specific copy of $\tilde{H}$, call it S and define a function g on $Z_2^{E \setminus |S|}$:

$$g(R) = f(R \cup S);$$

in other words g is the same as f but its domain is all CSP's with set of constraints disjoint from that of S. By analyzing g we will show that $H \setminus \tilde{H}$ is balanced, and hence H is balanced.

Remark 23 In order to prove that $H \setminus \tilde{H}$ is balanced and conclude that H is balanced, we must interpret correctly the meaning of average degree, orbit, expectation, etc. in our new space. For any CSP in new space the constraint set is well defined, we define the vertex set to be only those vertices not in $v(S)$, and the new orbit is defined by all automorphisms of the complete CSP leaving S fixed.

As we have seen in (11)

$$\hat{g}(R) = \sum_{S' \subseteq S} \hat{f}(R \cup S') U_{S'}(S). \tag{18}$$

Since the weight of $\hat{f}^2$ on any orbit of size $|\Theta(G)|$ is no more than a constant ($||f||_2 < 1$) we have that $|\hat{f}(G)| \leq c/\sqrt{|\Theta|}$ (and in case of H we have equality for some c). Recalling that $\mathbf{Ex}[S'] = |\Theta(S')| \prod_{e \in S'} p_e$ and $U_{S'}(S) = -\prod_{e \in S'} \sqrt{q_e/p_e}$ this gives

$$|\hat{f}(R \cup S') U_{S'}(S)| \leq c/\sqrt{\mathbf{Ex}[S']}.$$

The fact that S is strictly balanced and of expectation bounded by a constant means that $\mathbf{Ex}[S'] \gg \mathbf{Ex}[S]$ for all proper subproblems S' of S.

So putting $R = H \setminus S$ in (18) we have that the dominant summand is the one corresponding to $S' = S$, thus

$$\hat{g}(H \setminus S) \sim \frac{|\hat{f}(H)|}{\sqrt{\prod_{e \in S} p_e}}. \tag{19}$$

Let $W = \sum_{G \in \tilde{\Theta}} \hat{g}^2(G)$ where $\tilde{\Theta}$ is the orbit of $H \setminus S$ in our new space. Note that

$$|\tilde{\Theta}(H \setminus \tilde{H})| = c' |\Theta(H)|.$$

Now (19) gives

$$W = c' |\tilde{\Theta}(H \setminus \tilde{H})| \hat{f}^2(H) / \left(\prod_{e \in S} p_e\right) = c / \left(\Theta(\tilde{H}) \left(\prod_{e \in S} p_e\right)\right) = c\mathbf{Ex}[\tilde{H}],$$

i.e. W is approximately constant.

The following computation shows that $pI(g)$ is no more than polylogarithmic:

We will now use Russo's lemma via the interpretation that pI is the expected number of "pivotal" constraints, to lend a term from percolation theory.

For a given monotone family of CSP's A a given $G \in A$ (which has n variables) a *pivotal* constraint is a constraint $e \in G$ such that $G \setminus e$ does not belong to A. The number of the pivotal constraints of $\widehat{CSP}_{n,p}(\mathcal{P})$ is a random variable C, and by Lemma 14, $\mathbf{Ex}[C] = pI$. Let $X_{\tilde{H}}$ be a random variable that counts the number of copies of $\tilde{H}$ in $\widehat{CSP}_{n,p}(\mathcal{P})$. By abuse of notation let S also be the event that "S appears in $\widehat{CSP}_{n,p}(\mathcal{P})$". (Recall S is a specific copy of $\tilde{H}$.) We have

$$\begin{aligned}\mathbf{Ex}[C|S] &= \sum \mathbf{Ex}[C|X_{\tilde{H}} = i] \Pr[X_{\tilde{H}} = i|S] \\ &= \sum \mathbf{Ex}[C|X_{\tilde{H}} = i] \Pr[X_{\tilde{H}} = i] \Pr[S|X_{\tilde{H}} = i] / \Pr[S] \\ &= \sum \mathbf{Ex}[C|X_{\tilde{H}} = i] \Pr[X_{\tilde{H}} = i].i / \mathbf{Ex}[\tilde{H}]\end{aligned}$$

The following claim verifies that this expression cannot differ by more than a polylogarithmic factor from the expression for the unconditioned expectation of C.

Claim: Let H be a strictly balanced CSP, and let p be such that $\mathbf{Ex}[H] < C$. Then for every r the probability of having more than $\log(n)$ copies of H in $\widehat{CSP}_{n,p}(\mathcal{P})$ is asymptotically less than n^{-k}.

Proof. Let B be a large integer to be determined, and divide the event of having $\log(n)$ copies of H into two subevents:

1. The number of copies of G that intersect other copies is at most B.

2. More than B copies intersect other copies.

Denoting $\log(n) - B = R$, the probability of 1 is bounded by a constant times $C^R/R!$, asymptotically less than $n^{-r}/2$ for any fixed B. This follows from the usual way if computing the moment of the random variable X_H.

The event 2 can be described by the existence of a subproblem a list of subproblems whose length is a function of B, $l(B)$. Any CSP in this list can be described by a union of at least B but no more than $2B$ copies of H. It follows from the fact that H is strictly balanced that for sufficiently large B the expected number of copies of any CSP in this list is smaller than n^{-r+1}, for all n large enough. So for sufficiently large B the probability of 2 is asymptotically less than $n^{-r}/2$. ∎

We now can proceed to analyze $H \setminus \tilde{H}$ in the same manner we analyzed H: The weight of $\hat{g}^2$ on the orbit of $H \setminus \tilde{H}$ is a constant $pI(g)$ is no more than a power of $\log(n)$, so we once again can divide into three cases as before. Even though pI is now logarithmic rather than constant our analysis is the same because once we know that $\mathbf{Ex}[\tilde{H}]$ is constant any CSP with expectation bounded by a power $\log(n)$ must also have bounded expectation. So we once again divide into 3 cases. In case 1 or 2 we are done, and in case 3 we iterate the computation again since the size of H was bounded a priori, this process will terminate after a finite number of steps. This shows that H can be built by taking $\tilde{H}$, adding a CSP which is the minimal strictly balanced subproblem of $H \setminus \tilde{H}$, and so on. This must result with a balanced CSP, since in each step we add a balanced CSP from our new spaces. We also saw that $\mathbf{Ex}[H]$ is also bounded this completes the proof of the lemma. ∎

To summarize of $\sum_{R \sim H} \hat{f}^2(R)$ is not too small, H must be modest:

1. $|H|$ is bounded by Lemma 15.

2. $\mathbf{Ex}[H]$ is bounded from below by (10).

3. H is balanced and with expectation bounded from above by Lemma 22.

7.1.4 Main Proof

Let L, c_1, and c_2 be fixed. Remember that H is modest if

- $|H| \leq L$.
- $c_1 \leq \mathbf{Ex}[H] \leq c_2$.
- H is balanced.

Let $S_1, \ldots, S_l$ be a list of all the modest CSP's. For any CSP S, let C_S be the set of all CSP's T on n variables such that the union of all copies of the S_i's that appear as a subproblem of T is isomorphic to S. We will subdivide the space of subproblems of the complete CSP on n variables into these disjoint sets and approximate f separately on each. While doing this we will define certain parts of our space in which rare events occur. Roughly speaking:

1. $\mathcal{C}_1$ is the union of C_S for S which are large.

2. $\mathcal{C}_2$ is the union of C_S for which $\mu(C_S)$ is small.

3. $\mathcal{C}_3$ is the union of C_S for which $\Pr[f = 1]$ is not close to 0 or 1.

Let $\mathcal{C}_\triangle$ be the union of the remaining sets. We will now define a sequence of approximations of f, according to small positive constants $\epsilon_1, \ldots, \epsilon_5$. Our final approximation of f will be equal to 1 on a CSP R if and only if R has a subproblem such that $C_S \subset \mathcal{C}$ and f is equal to 1 on most of C_R.

Step 1: Let $g_1 = \sum \hat{f}(S) V_S$ where the sum is only on modest CSP's S, and $V_S = \sum_{G \in \Theta(S)} U_G$. By Lemma 20, we can choose the bounds for modest CSP's so that $||f - g||_2 \leq \epsilon_1$.

Step 2: Let $\mathcal{C}_1$ be the union of all C_S such that $X_{S_i}(S) \geq l\mathbf{Ex}[S_i]/\epsilon_2$, for some i, i.e. the number of copies of S_i appearing is much more than the expected. The measure of $\mathcal{C}_1$ is no more than ϵ_2. Define g_2 to be equal identically to 1 on C_1, and equal to g_1 otherwise. So $||f - g_2||_2 \leq \epsilon_1 + \sqrt{\epsilon_2}$.

Step 3: Recalling that $\mathbf{Ex}[S_i]$ is bounded we have that the number of CSP's in

$$\{S : C_S \neq \emptyset, C_S \not\subset C_1\}$$

is bounded, i.e. we have a bound on the number of subsets C_S on which g_2 is not identically 1. Say we have M such sets. Let $\mathcal{C}_2$ be the union of all C_S not in $\mathcal{C}_1$ such that $\mu(C_S) < \epsilon/M$. So $\mu(\mathcal{C}_2) \leq \epsilon_3$.

The reason we treat this part of our space separately is because we need a lower bound on the measure of C_G in the proof of Lemma 29. Let $\mathcal{C}$ be the union of all remaining sets C_S, those not in $\mathcal{C}_1 \cup \mathcal{C}_2$.

Remark 24 Note that if S is such that $C_S \subset \mathcal{C}$, then S must be balanced and the average degree of the constraint hypergraph of S must be the same as all modest CSP's. Furthermore the size of S is bounded from above, and $\mathbf{Ex}[S]$ is bounded from above and below.

For any $C_S \subset \mathcal{C}$ and for any CSP $H \in C_S$ define

$$g_3(H) = \mathbf{Ex}[g_2 | C_S],$$

i.e. we replace g_2 by its conditional expectation in C_S. Define g_3 to be equal to 1 on all CSP's not in $\mathcal{C}$. We will show that for any $C_S \subset \mathcal{C}$, g_3 is close to g_2 because g_2 is almost constant on C_S in the sense that for any constant δ

$$(\{g_2(T) - \mathbf{Ex}[g_2 | C_S] > \delta\} | T \in C_S) \to 0 \tag{20}$$

This will be proven in Lemma 29 in the following subsection. Recalling there are only M sets C_S in $\mathcal{C}$, we get that by choosing δ small enough by proper choice of ϵ_1, ϵ_2, ϵ_3 we have

$$||g_3 - f|| \leq \epsilon_4/2$$

Step 4: For any S let g_4 on the CSP's in C_S be identically 0 or 1 according to which approximate f better.

We now wish to compare two approximations of f: g_3 which is constant on each C_S, and g_4 which is the best approximation that is both constant on each C_S and Boolean. Let h be the best possible L^2 approximation f that is constant on C_S. A simple calculation shows that

$$h|_{C_S} = \Pr[f(R) = 1 | R \in C_S].$$

The following inequality follows by summing over each separately.

$$||f - g_4||_2^2 \leq 2||f - h||_2^2 \leq 2||f - g_3||_2^2 = \epsilon_4. \tag{21}$$

g_4 is the characteristic function of a family B that is a candidate to be the family guaranteed by the theorem. Yet we do not know that B is monotone and that $||B||$ is small.

Step 5: We now define g_5, a monotone Boolean function which is constant on each C_S. We will define it such that if $R \subset S$, then $g_5|C_R \leq g_5|C_S$. call a CSP S *decisive* if

$$\Pr[f(T) = 1 | T \in C_S] \notin (\sqrt{\epsilon_4}, 1 - \sqrt{\epsilon_4}).$$

Let $\mathcal{C}_3$ be the union of all sets C_S for S which is not decisive. From (21) we have that $\mu(\mathcal{C}_3) \leq \sqrt{\epsilon_4}$. Let $\mathcal{C}_4$ be $\mathcal{C} \setminus \mathcal{C}_3$. For any S define g_5 on C_S to be equal to 1 if and only if S has a subproblem R such that $C_R \subset \mathcal{C}_4$, with g_4 equal to 1 on C_R. Since the union $\mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3$ is of small measure, the alterations on C_S that belong to these parts of our space do not affect our approximation much. We will show in Lemma 25 below that if $R \subset S$, and C_R and C_S belong to $\mathcal{C}_4$, then

$$\mathbf{Ex}[f|C_R] \geq 1 - \sqrt{\epsilon_4} \Rightarrow \mathbf{Ex}[f|C_s] \geq 1 - 2\sqrt{\epsilon_4}$$

and hence on the sets $C_S \subset \mathcal{C}_4$ we do not alter our approximation at all. Therefore choosing sufficiently small ϵ_4 we have

$$||g_5 - f||_2^2 \leq \epsilon.$$

g_5 is the characteristic function of a symmetric monotone family B with minimal CSP's which are balanced and of bounded size, and

$$\mu(B \Delta A) \leq \epsilon.$$

This completes the proof of Theorem 11. Furthermore by choosing proper choices of ϵ_i, $i = 1, \ldots, 4$, we can define $\mathcal{C}_4$ such that, for $C_R \subset \mathcal{C}_4$, $\mathbf{Ex}[f|C_R] > 1 - \epsilon$. Since we have a bound on $\mathbf{Ex}[R]$ and $|R|$ for all CSP's R, any one of them is a candidate to be the CSP guaranteed by Theorem 12. Let r be a specific copy of such a CSP R, and let B_r be the space of all CSP's having r as a subproblem. Let $C_r = B_r \cap C_R$. From symmetry

$$\mathbf{Ex}[f|C_r] = \mathbf{Ex}[f|C_R]$$

and from positive correlation of increasing events

$$\mathbf{Ex}[f|B_r] \geq \mathbf{Ex}[f|C_r].$$

In other words, conditioning on the appearance of r the expectation of f is at least $1 - \epsilon$. This concludes the proof of Theorem 12.

Lemma 25 *Let $R \subset T$ be CSP's, and let $C_R, C_T, \mathcal{C}_4$ be defined as above. Suppose $C_R \subset \mathcal{C}_4$ and $C_T \subset \mathcal{C}_4$. Then*

$$\mathbf{Ex}[f|C_T] > \mathbf{Ex}[f|C_R] - o(1)$$

Proof. Recall that all the four quantities $\mathbf{Ex}[R], \mathbf{Ex}[T], \mu(C_R), \mu(C_T)$ were, by definition of $\mathcal{C}_\triangle$ and by Remark 24 bounded from above and from below by some constants (that depend on ϵ) and that R and T must be balanced, with a bounded size and average degree δ equal to that of the modest CSP's.

Let $r \subset t$ be specific copies of R and T. Define B_r to be the space of all CSP's that have e as a subproblem. We will consider all subspaces with the induced conditional probability. Let

$$C_r = B_r \cap C_R,$$

and define C_t analogously. Note that C_R is the disjoin union of $\Theta(R)$ sets isomorphic to C_r. By symmetry we have:

$$\mathbf{Ex}[f|C_R] = \mathbf{Ex}[f|C_r]$$

and

$$\mathbf{Ex}[f|C_T] = \mathbf{Ex}[f|C_t].$$

We will define a mapping $\sigma : C_t \to c_r$ such that σ is one to one and measure preserving with respect to the conditional measure on C_t and its image. More precisely, if μ_1, μ_2 are respectively the conditional measure on C_t and $\sigma(C_t)$, then for any $G \in C_t$, $\mu_1(G) = \mu_2(\sigma(G))$. Furthermore σ will be such that

$$f(\sigma(G)) = 1 \Rightarrow f(G) = 1$$

and hence

$$\mathbf{Ex}[f|\sigma(C_t)] \geq \mathbf{Ex}[f|C_r] - o(1).$$

Therefore it will suffice to show that

$$\mathbf{Ex}[f|\sigma(C_t)] \geq \mathbf{Ex}[f|C_r] - o(1). \tag{22}$$

The definition of σ is very simple: For every $G \in C_t$ define $\sigma(G)$ to be the CSP obtained from G by deleting the constraints in $t \setminus r$. Obviously $\sigma(G) \in C_r$. The question is which CSP's in C_r are not in the image of σ. Let $W_1 = \sigma(C_t)$ and $W_2 = c_r \setminus W_1$. The CSP's in W_2 can be classified into two types:

1. CSP's that have a constraint in $t \setminus r$.
2. CSP's G such that $G \cup t \notin C_t$.

Roughly, the reason for a CSP G to belong to W_2 is that there exist a subproblem s in G such that $s \cup t$. is a union of modest CSP's, whereas $s \cup r$ is not.

For a CSP $G \in C_r$ define

$$\alpha(G) = \Pr[\pi(G) \in W_2] \tag{23}$$

where π is a permutation of the vertices leaving the vertices of r fixed, chosen uniformly at random from such permutations. Let

$$\alpha = \frac{\mu(W_2)}{\mu(C_r)}.$$

Let μ_r be the conditional measure on C_r. So

$$a = \int C_r \alpha d_{\mu_r}.$$

On the other hand,

$$1-a=\mu(\alpha(C_t))/\mu(C_r)=\mu(C_t)/(\prod_{e\in T\setminus R}p_e\mu(C_r))$$

$$=\frac{\mu(C_T)}{|\Theta(T)|}\frac{|\Theta(R)|}{\mu(C_R)}\frac{\prod_{e\in R}p_e}{\prod_{e\in T}p_e}=\frac{\mu(C_t)\mathbf{Ex}[R]}{\mu(C_R)\mathbf{Ex}[T]}.$$

Recalling the properties of R and T this shows a is bounded away from 1. Using this and the fact that f is constant on isomorphism classes we have

$$\mathbf{Ex}[f|W_1]=\frac{\int_{C_r}f.(1-\alpha)d_{\mu_r}}{(1-a)}.$$

But lemma 26 below shows that α is essentially constant (=a) on C_r, and 22 follows. ■

Lemma 26 *Let $\alpha=\alpha(G)$ be as defined in (23), where G is chosen at random from C_r by the measure μ_r. Then*

$$\mathbf{Var}[\alpha]=o(1).$$

Proof. First, we would like to shift to working in B_r, a space with a convenient product measure, and to this end we will extend the definition of $\alpha(G)$ to all CSP's $G\in B_r$. Note that a CSP in B_r almost surely has no constraints in $t\setminus r$, and we will disregard the exceptions to this rule in our calculations. For any CSP G let $V(G)$ denote the set of variables of G. Let $k=|V(t)|-|V(r)|$. Order the variables in K, $v_1,\dots,v_k$, For any ordered set $x=\{x_1,\dots,x_k\}$ outside of $V(r)$, let π_x be a permutation of the variables of the complete CSP leaving $V(r)$ fixed such that $\pi(x_i)=v_i$ for $i=1,\dots,k$. For a given CSP $G\in C_r$ define a set of variables x to be *problematic* if

$$\pi_x(G)\le W_2.$$

This does not depend on the choice of π. Note that $\alpha(G)$ is exactly the proportion of problematic k-sets. For any problematic set x, we can find a set of constraints and variables y in $G\setminus(r\cup x)$ that "are a reason" for x being problematic. More precisely $\pi_x(y)\cup t$ is a union of modest CSP's. (In our example these are the paths of length 3 added between two variables in t.) Call such a set an extension of x. We now want to define extensions and problematic sets for any CSP in B_r and a set of variables and constraints y which is disjoint from r and x, we say that y is an extension of x if

$$(\pi_x(y\cup x)\cup r)\in C_r,$$

but

$$(\pi_x(y\cup x)\cup t)\notin C_t.$$

Remark 27 Note that we view y as a set of constraints and variables and not as a CSP, indeed some constraints in y may be such that their variables are not in y. We do require, however, that $y\cup x$ be a CSP.

Now extend the definition of a problematic set of variables in a CSP in B_r to include any set of vertices that has an extension. For every $G\in B_r$, define $\alpha(G)$ to be the proportion of the problematic sets among all sets of size k. This definition of α coincides with the previous one on C_r.

Now before shifting to work on B_r note that

$$\mu(C_r)/\mu(B_r) = \mu(C_R)/\left(|\Theta(R)|\prod_{e\in R} p_e\right) = \mu(C_R)/\mathbf{Ex}[R].$$

From the definition of $\mathcal{C}_4$, and the fact that $C_r \subset \mathcal{C}_4$, $\mu(C_R)$ is bounded from below and $\mathbf{Ex}[R]$ is bounded from above (by bounds that depend on ϵ), and hence the relative measure of C_r in B_r is non-negligible, so

$$\mathbf{Var}[\alpha|B_r] = o(1) \Rightarrow \mathbf{Var}[\alpha|C_r] = o(1).$$

Lemma 28 below asserts that $\mathbf{Var}[\alpha|B_r] = o(1)$, hence our result follows. ■

Lemma 28

$$\mathbf{Var}[\alpha|B_r] = o(1).$$

Proof. Define $X = n_k\alpha$, where $n_k = n!/(n-k)!$. $X(G)$ os the number of problematic k-sets in G. Let $x_1, \ldots, x_{n_k}$ be the indicator random variables of the events of the corresponding sets being problematic. So $X = \sum x_i$.

If $\mathbf{Ex}[X] = o(n^k)$, then $\mathbf{Ex}[\alpha] = o(1)$ and since $0 \le \alpha \le 1$, $\mathbf{Var}[\alpha] = o(1)$. Hence we may assume that $\mathbf{Ex}[X] = \Omega(n^k)$ and strive to prove that $\mathbf{Var}[X] = o(n^{2k})$. We have

$$\mathbf{Var}[X] = \sum_i\sum_j(\mathbf{Ex}[x_ix_j] - \mathbf{Ex}[x_i]\mathbf{Ex}[x_j]).$$

For $i \neq j$, let $x_i \circ x_j$ be the random variable indicating the event that there exist *constraint disjoint* extensions of the corresponding sets. The BK inequality [32] implies

$$\mathbf{Ex}[x_ix_j] \le \mathbf{Ex}[x_i]\mathbf{Ex}[x_j].$$

Let $x_i \diamond x_j = x_ix_j - x_i \circ x_j$. We have

$$\mathbf{Var}[X]\sum \mathbf{Var}[x_i] + \sum_{i\neq j}\mathbf{Ex}[x_i \diamond x_j] = \sum_{i\neq j}\mathbf{Ex}[x_i \diamond x_j] + o(n^{2k}). \tag{24}$$

We now need some notation regarding CSP's in space B_r. Define two CSP's G and G' to be of the same isomorphism type if there exists a permutation of the vertices not in $V(r)$ that takes G to G'. Define $\tilde{\mathbf{Ex}}[G]$ as the expected number of copies isomorphic to G in a random CSP in B_r.

Let the average degree of the modest CSP's be δ. For any set of constraints and variables define the average degree to be the ratio between the number of constraints in B_r. So, for an extension y of a set x the average degree is the ratio between the number constraints in y and the number of vertices of y (those not in x and not in $V(r)$). (Recall that y is not necessarily a proper CSP in the sense that not all constraints in y are between vertices that belong to y.) Recall that $t \cup \pi_x(y)$ is the union of modest CSP's, and hence has average degree at least δ. Hence the average degree of y must also be at least δ.

We will say an extension y is *nice* if

1. It is a minimal extension.

2. Its average degree is δ.

3. For any $z \subset y$ such that $z \cup t$ is a CSP the average degree of z is no larger than δ.

The reason for defining this notion is that our calculations are much simpler when considering such extensions. Whenever $x_i \diamond x_j = 1$, there are minimal extensions causing this. Furthermore, for a given set x, the probability if having a minimal extensions with average degree of it, or any subextension larger than δ, is $o(1)$. Hence we may concentrate on the events caused by nice extensions: Let $x_i \star x_j$ be the indicator random variable if the event indicated by $x_i \diamond x_j$, where there exist nice extensions causing this event. We have

$$\sum \mathbf{Ex}[x_i \diamond x_j - x_i \star x_j] = o(n^{2k}),$$

and hence it suffices to show

$$\mathbf{Ex}[\sum x_i \star x_j] = o(n^{2k}) \tag{25}$$

For an extension y let $\mathrm{Cl}(y)$ denote the CSP's whose constraints are the constraints in y. We will need the following property of nice extensions:

Claim: Let x be a set of variables and y a nice extension. Any subproblem of $\mathrm{Cl}(y) \setminus r$ has average degree smaller than δ.

(Note that the claim deals with actual CSP's and not extensions, i.e. all constraints come with their variables.)

Proof of Claim: Let y be a nice extension of x and assume for simplicity if notation that $x = V(t) \setminus V(r)$. Note that from minimality of y there exists a modest CSP S such that $\mathrm{Cl}(y) \subseteq S \subseteq y \cup t$. This S is the disjoint union of three sets:

1. $S \cap r$.

2. $S \cap (t \setminus r)$.

3. All the rest, namely $y \setminus t$.

Note that Parts 2 and 3 are not necessarily proper CSP's, i.e. we may have constraints with less than k variables belonging to them.

Part 3 does not have average degree less than δ because y is nice. Part 2 cannot have average degree larger than δ, or else its union with r also have a larger average degree, but this union is a subproblem of t which has no such subproblems. Hence Part 1 cannot have average degree *smaller* than δ. As a subproblem of r it cannot have larger degree either, and hence has average degree exactly δ. Now if $z \subset (\mathrm{Cl}(y) \setminus r)$ has average degree δ, then $z \cup (S \cap r)$ is modest (it is a subproblem of S which is modest, and has the correct average degree). Hence $r \cup z$ is a union of modest CSP's, which is a contradiction: since y is an extension, there exists a CSP in C_r with z as a subproblem, but in C_r the union of all modest CSP's is r.

Let y be a nice extension of a set x. The CSP's in B_r consisting of the variables of x and constraints of y can take on a finite number of isomorphism types $G_1, \ldots, G_d$. For all these CSP's we have

$$\tilde{\mathbf{Ex}}(G_i) = O(n^k). \tag{26}$$

This follows from the fact that there is a copy of G_i, say g, whose union with t is a union of modest CSP's. We have $\mathbf{Ex}[g \cup t] < c$. But

$$\mathbf{Ex}[g \cup t] \geq c\mathbf{Ex}[t]\tilde{\mathbf{Ex}}(g)/n^k$$

and $\mathbf{Ex}[t]$ is bounded from below.

We now can prove 25:

When summing $\mathbf{Ex}[x_i \star x_j]$ we use the fact that if an extension of type R_i intersects an extension of type R_j and their intersection is of type H they form an event such that the expected number of isomorphic events is $\tilde{\mathbf{Ex}}[R_i]\tilde{\mathbf{Ex}}[R_j]/\tilde{\mathbf{Ex}}[H]$. But from the fact that R_i is nice it follows from Claim that the average degree of H is smaller than δ, or in other words $\tilde{H} \to \infty$, and

$$\tilde{\mathbf{Ex}}[R_i]\tilde{\mathbf{Ex}}[R_j]/\tilde{\mathbf{Ex}}[H] = o(n^{2k}).$$

Since we have a finite number of such contributions this gives the desired bound. ■

The last brick missing in the proof of the Theorems 11 and 12 is the following lemma. We defined g_2, an approximation of f and $\mathcal{C}$ as the union of all sets C_G with the following properties:

- $|G|$ was bounded from above.
- $\mathbf{Ex}[G]$ was bounded from above and below.
- $\mu(C_G)$ was not too small.

We promised to show that g_2 is almost determined by the number of appearances of the CSP's S_i, in the sense that if $C_S \subset \mathcal{C}$, then for any constant $\delta > 0$

$$\Pr[\{|g_2(T) - \mathbf{Ex}[g_2|C_S] > \delta\}|T \in C_S] \to 0. \tag{27}$$

Recalling the Fourier expansion of g_2 it is sufficient to show this for the function $\hat{f}(S)V_S$, where S is modest, or using (9):

Lemma 29 *Let G be such that $C_G \subset \mathcal{C}$. Let S be modest and $V = V_S$. Then:*

$$\forall \delta > 0 \Pr[\{|V(T) - \mathbf{Ex}[V|C_G]| > \sqrt{p}^{-|S|}\delta\}|T \in C_G] \to 0.$$

Proof. Note that all the modest CSP's have the same average degree. Furthermore $\mathbf{Ex}[S]/\mathbf{Ex}[R]$ is bounded for any R that is a subproblem of a modest CSP S. Recalling Lemma 16 we wish to show that $\frac{X_R}{\mathbf{Ex}[R]}$ is almost constant on each C_G for $R \subseteq S$. So let us calculate the conditioned variance of X_R in a given C_G. If R is modest, then X_R is constant. So we may concentrate on R which is a subproblem of one of the S_i, but not modest itself, i.e. $\mathbf{Ex}[R]$ is large. Recall that Lemma 19 gave the following expression for the non-conditioned variance of X_R:

$$\mathbf{Var}(X_R) \asymp \mathbf{Ex}[R]^2(\sum 1/\mathbf{Ex}[H])$$

where the sum is over all non-empty subproblems of R. If none of the modest CSP's are subproblems of R this is $o(\mathbf{Ex}[R]^2)$, since for every subproblem H, $\mathbf{Ex}[H]$ is also large. So the standard deviation of X_R in this case is $o(\mathbf{Ex}[R])$. Obviously conditioning on an event whose probability is bounded away from zero (being in C_G) cannot change this. Let us consider then, the variance of X_R when some S_i is a subproblem of R. C_G was defined b the fact that the union of the modest CSP's appearing was isomorphic to G. Let g be a *specific* copy of G. Let B_g be the space consisting of all CSP's that have g as a subproblem with the probability measure induced by the conditional probability in $\widehat{CSP}_{n,p}(\mathcal{P})$. As in Remark 23 we define CSPs, orbits, expected number copies of CSP's, etc. in our space in natural way. Let $c_g = C_G \cap B_g$. From symmetry we ger that the expectation and variance of X_R conditioned on being in C_G is the same as conditioning on being C_g. Focusing our attention on C_g every copy of R must have g as a subproblem. Therefore X_R now depends on the appearance of copies of certain CSP's $T_1, \ldots, T_k$, such that $T_i \cup g \sim R$. So we may now define X_T so that $X_R = X_T$ (in the space C_g), but X_T counts the appearance of copies of the T_i's. Since g is the union of all modest CSP's in any CSP in C_g we may assume T_i has no modest subproblem in B_g, and $\mathbf{Ex}[H] \to \infty$ for all $H \subset T_i$.

Now, a simple calculation shows that

$$\mu(C_G \cap B_g)/\mu(B_g) = \mu(C_G)/\mathbf{Ex}[G].$$

and from the definition of $\mathcal{C}$, $\mathbf{Ex}[G]$ is bounded from above and $\mu(C_G)$ from below, hence if $\mathbf{Var}[X_T] = o(\mathbf{Ex}[X_T]^2)$ conditioned on being in B_g the same must be true on C_g.

We now present the calculation done in Lemma 19 in the same manner as done in the proof of Lemma 28 with the sum

$$\sum \mathbf{Ex}[XY] - \mathbf{Ex}[X]\mathbf{Ex}[Y]$$

and the expectations as defined in the space B_g. We get that

$$\mathbf{Var}[X_T] \leq c\mathbf{Ex}[X_T]^2 / \min_{H \subset T_i} \mathbf{Ex}[H].$$

We conclude that this is $o(\mathbf{Ex}[T]^2)$, so the standard deviation of X_T is indeed $o(\mathbf{Ex}[R])$.

The above consideration show that $X_R/\mathbf{Ex}[R]$ is almost a constant on any set C_G (its standard deviation is $o(1)$), and hence by Lemma 16 $\hat{f}(S)V_S = cV_S/\sqrt{\Theta(S)}$ is indeed almost constant on our subsets. Moreover for S such that $\sum_{H \in \Theta(S)} \hat{f}^2(G) = c$

$$\hat{f}(S)V_S = V_S\sqrt{c/\Theta(S)} \sim \sqrt{c/\mathbf{Ex}[S]} \left(\sum (-1)^{|S|-|R|} \frac{\tilde{\mathbf{Ex}}[R]}{\mathbf{Ex}[R]} \right)$$

where $\tilde{\mathbf{Ex}}$ is the conditioned expectation of X_R. This completes the proof of the lemma. ∎