Why almost all k-CNF formulas are easy

Amin Coja-Oghlan[1], Michael Krivelevich[2] and Dan Vilenchik[3]

[1] Institute for Informatics, Humboldt-University, Berlin, Germany.
coja@informatik.hu-berlin.de.
[2] School of Mathematical Sciences, Tel-Aviv University, Tel-Aviv, Israel.
krivelev@post.tau.ac.il.
[3] School of Computer Science, Tel-Aviv University, Tel-Aviv, Israel.
vilenchi@post.tau.ac.il.

Abstract. Finding a satisfying assignment for a k-CNF formula ($k \geq 3$), assuming such exists, is a notoriously hard problem. In this work we consider the uniform distribution over satisfiable k-CNF formulas with a linear number of clauses (clause-variable ratio greater than some constant). We rigorously analyze the structure of the space of satisfying assignments of a random formula in that distribution, showing that basically all satisfying assignments are clustered in one cluster, and agree on all but a small, though constant proportion, number of variables. This observation enables us to describe a polynomial time algorithm that finds *whp* a satisfying assignment for such formulas, thus asserting that most satisfiable k-CNF formulas are easy (whenever the clause-variable ratio is greater than some constant). This should be contrasted with the setting of very sparse k-CNF formulas (which are satisfiable *whp*), where experimental results show some regime of clause density to be difficult for many SAT heuristics. One explanation for this phenomena, backed up by partially non-rigorous analytical tools from statistical physics, is the complicated clustering of the solution space at that regime, unlike the more "regular" structure that denser formulas possess. Thus in some sense, our result rigorously supports this explanation.

key words: computational and structural complexity, algorithms and data structures, message passing algorithms, SAT.

1 Introduction

Constraint satisfaction problems play an important role in many areas of computer science, e.g. computational complexity theory [9], coding theory [16], and artificial intelligence [26], to mention just a few. The main challenge is to devise efficient algorithms for finding satisfying assignments (when such exist), or conversely to provide a certificate of unsatisfiability. One of the best known examples of a constraint satisfaction problem is k-SAT, which is the first to be proven as NP-complete. Although satisfactory approximation algorithms are known for several NP-hard problems, the problem of finding a satisfying assignment (if such exists) is not amongst them. In fact, Håstad [17] proved that it is NP-hard to approximate MAX-3SAT (the problem of finding an assignment that satisfies as many clauses as possible) within a ratio better than 7/8.

In trying to understand the inherent hardness of the problem, many researchers analyzed structural properties of formulas drawn from different distributions. One such natural distribution is the following: fix $c, n > 0$ (c may depend on n), choose $m = cn$ clauses uniformly at random out of $8\binom{n}{3}$ possible ones. We denote this distribution by $\mathcal{P}_{n,m}$. Despite its simplicity, many essential properties of this model are yet to be understood. In particular, the hardness of deciding if a random formula is satisfiable, and finding a satisfying assignment for a random formula, are both major open problems [10, 23].

1.1 Our Contribution

Remarkable phenomena occurring in the random model $\mathcal{P}_{n,m}$ are **phase transitions**. With respect to the property of being satisfiable, such a phase transition takes place too. More precisely, there exists a threshold $d_k = d_k(n)$ such that a k-CNF formula with clause-variable ratio greater than d_k is not satisfiable *whp*[4], while one with ratio smaller than d_k is [15]. In this work we consider satisfiable k-CNF formulas with cn clauses, c greater than some sufficiently large constant. In this regime almost all formulas are not satisfiable, and therefore we consider the following natural extension of $\mathcal{P}_{n,m}$, which we denote by $\mathcal{P}^{\text{sat}}_{n,m}$: fix $c, n > 0$ (c may depend on n), choose $m = cn$ clauses uniformly at random out of $8\binom{n}{3}$ possible ones, *conditioned* on the fact that the received formula is satisfiable. To simplify the presentation we consider the most popular setting – the case $k = 3$, namely random 3SAT, and remark that our results extend to any fixed k.

Our contribution is composed of three parts. The **first** part *rigorously* establishes the following fact:

Theorem 1. *There exists a polynomial time algorithm that* whp *finds a satisfying assignment for 3CNF instances from $\mathcal{P}^{\text{sat}}_{n,m}$, $m \geq C_0 n$, C_0 a sufficiently large constant.*

Thus we partially answer the open problem asking to decide the hardness of $\mathcal{P}_{n,m}$ [10, 23]. Specifically, we assert that for all but a vanishing fraction of satisfiable 3CNF formulas over n variables with m clauses, one can efficiently find a satisfying assignment (whenever m/n is greater than some constant). Our proof of Theorem 1 is constructive – that is, we present an algorithm that meets the requirements of Theorem 1.

The **second** part of our result concerns another exciting area. One of the most surprising recent developments in satisfiability problems comes from statistical physics. More specifically, in their well-known work, Mezard, Parisi and Zecchina [6] designed a new algorithm, known as Survey Propagation, for solving k-SAT instances. A particularly dramatic feature of this method is that it appears to remain effective in solving very large instances of random k-SAT even with densities very close to the conjectured satisfiability threshold, a regime where other algorithms (e.g., the WalkSAT method [27]) typically fail. Nonetheless, despite the considerable progress to date, the reasons underlying the remarkable performance of Survey Propagation are not yet fully understood, let alone a rigorously analyzed.

The difficulty that Survey Propagation apparently overcomes lies in the complicated structure of the solution space of such formulas. That is, the *conjectured* picture, some supporting evidence of which were proved rigorously for $k \geq 8$ [24, 1, 25], is that typically random k-CNF formulas in the near-threshold regime have an exponential number of **clusters** of satisfying assignments. While any two assignments in distinct clusters disagree on at least εn variables, any two assignments within one cluster coincide on $(1-\varepsilon)n$ variables. Furthermore, each cluster has a linear number of **frozen** variables (a variable is said to be *frozen* in some cluster if *all* satisfying assignments within that cluster assign it in the same way). The algorithmic difficulty with such a clustered solution space seems to be that most known algorithms do not "steer" into one cluster but try to find a "compromise" between the satisfying assignments in distinct clusters, which actually is impossible.

Complementing this picture *rigorously*, we show that typically for satisfiable 3CNF formulas in the above-threshold regime the solution space contains only one cluster, though its size may be exponential in n. Formally,

[4] Writing *whp* we mean with probability tending to 1 as n goes to infinity.

Theorem 2. *Let $\mathcal{F}$ be random 3CNF from $\mathcal{P}^{\text{sat}}_{n,m}$, $m \geq C_0 n$, C_0 a sufficiently large constant. Then* whp *$\mathcal{F}$ enjoys the following properties:*

1. *All but $e^{-\Theta(m/n)}n$ variables are frozen.*
2. *The formula induced by the non-frozen variables decomposes into connected components of at most logarithmic size.*
3. *Letting $\beta(\mathcal{F})$ be the number of satisfying assignments of $\mathcal{F}$, we have $\frac{1}{n}\log\beta(G) = e^{-\Theta(m/n)}$.*
4. *Any two satisfying assignments differ on at most $e^{-\Theta(m/n)}n$ variables*

Combining Theorems 1 and 2 supports the following common thesis: the main key to understanding the hardness (even experimental one) of a certain distribution over satisfiable formulas lies in the structure of the solution space of a typical formula in that distribution. Specifically, our results show (at least in our setting) that typically when a formula has a single cluster of satisfying assignments, though its volume might be exponential, then the problem is "easy". On the other hand, when the clustering is complicated, for example in the near threshold regime, experimental results predict that many "simple" heuristics fail, while "heavy machinery" such as Survey Propagation works. Heightening this last point, consider the recent work in [12], where the naïve Warning Propagation algorithm is rigorously shown to work *whp* for 3CNF formulas taken from a somewhat different distribution than the one we consider, nonetheless (as we shall prove) sharing with $\mathcal{P}^{\text{sat}}_{n,m}$ the same simple cluster structure. Fitting the result in [12] to our perspective – when the clustering is simple, then a simple message passing algorithm works (Warning Propagation), when the clustering is complicated, then only a much more complicated message passing algorithm is known (and even this only experimentally) to work (Survey Propagation).

The **third** part of our result is more "philosophical" in nature. As we already mentioned, the event of a random formula in $\mathcal{P}_{n,m}$ being satisfiable, when m/n is some constant above the satisfiability threshold, is very unlikely. Therefore, the distribution $\mathcal{P}^{\text{sat}}_{n,m}$ differs from the $\mathcal{P}_{n,m}$ distribution significantly. In effect, many techniques that have become standard in the study of random instances (3CNF formulas and grahps) just do not carry over to $\mathcal{P}^{\text{sat}}_{n,m}$ – at least not directly. In particular, the contriving event of being satisfiable causes the clauses in $\mathcal{P}^{\text{sat}}_{n,m}$ to be dependent.

The inherent difficulty of $\mathcal{P}^{\text{sat}}_{n,m}$ has led many researchers to consider the more approachable, but considerably less natural, **planted distribution**, pioneered by Kučera [22] in the context of graph coloring. In the planted distribution, which we denote by $\mathcal{P}^{\text{plant}}_{n,m}$, one first fixes some satisfying assignment, and then includes m clauses uniformly at random out of $7\binom{n}{3}$ clauses that are satisfied by it. This of course guarantees that the formula is satisfiable. Planted solution distributions are favored by many researchers in the context of SAT [14, 4, 20], but also for other graph optimization problems such as max clique, min bisection, and coloring [2, 3, 5, 19, 11], to mention just a few.

Of course the $\mathcal{P}^{\text{plant}}_{n,m}$ model is somewhat artificial and therefore provides a less natural model of random instances than $\mathcal{P}^{\text{sat}}_{n,m}$. Nevertheless, devising new ideas for analyzing $\mathcal{P}^{\text{sat}}_{n,m}$ we show that $\mathcal{P}^{\text{sat}}_{n,m}$ and $\mathcal{P}^{\text{plant}}_{n,m}$ actually share many structural properties such as the existence of a single cluster of solutions. As a consequence, we can prove that a certain algorithm, designed with $\mathcal{P}^{\text{plant}}_{n,m}$ in mind, works for $\mathcal{P}^{\text{sat}}_{n,m}$ as well (this algorithm is used to prove Theorem 1). In other words, by presenting new methods for analyzing heuristics on random instances, we can show that algorithmic techniques invented for the somewhat artificial planted model extend to the canonical uniform setting.

We proceed with related work and a detailed exposition of our techniques.

1.2 Related Work and Techniques

Almost all exact polynomial-time heuristics suggested so far for random instances (either SAT or graph optimization problems) were analyzed when the input is sampled according to a planted-solution distribution, or various semi-random variants thereof. Alon and Kahale [2] suggest a polynomial time algorithm based on spectral techniques that *whp* properly k-colors a random graph from the planted k-coloring distribution (the distribution of graphs generated by partitioning the n vertices into k equally-sized color classes, and including every edge connecting two different color classes with probability $p = p(n)$), for graphs with average degree greater than some constant. In the SAT context, Flaxman's algorithm, drawing on ideas from [2], solves *whp* planted 3SAT instances where the clause-variable ratio is greater than some constant. Also [13, 12, 21] address the planted 3SAT distribution.

On the other hand, very little work was done on non-planted distributions, such as $\mathcal{P}_{n,m}^{\text{sat}}$. In this context one can mention the work of Chen [7] which provides an *exponential* time algorithm for $\mathcal{P}_{n,m}^{\text{sat}}$ with m/n greater than some constant. Ben-Sasson et al. [4] also study $\mathcal{P}_{n,m}^{\text{sat}}$ but with $m/n = \Omega(\log n)$, a regime where $\mathcal{P}_{n,m}^{\text{sat}}$ and $\mathcal{P}_{n,m}^{\text{plant}}$ coincide (since typically there is only one satisfying assignment). [4] ask whether one can characterize $\mathcal{P}_{n,m}^{\text{sat}}$ for $m/n = o(\log n)$, and in particular they ask whether there exists a *polynomial* time algorithm that finds *whp* a satisfying assignment in this regime. In this work we answer their question positively. One should also mention the recent work of [8], where the uniform distribution over k-colorable graphs with average degree greater than some constant is analyzed. Specifically, [8] shows that a similar clustering phenomenon to the one described in Theorem 2 also occurs for k-colorable graphs with constant average degree. Furthermore, [8] shows that the algorithm by Alon and Kahale [2] works *whp* for such graphs as well. The techniques that we use are similar in flavor to the ones introduced in [8], though k-SAT is fundamentally different from k-colorability.

To obtain our results, we use two main techniques. As we mentioned, $\mathcal{P}_{n,m}^{\text{plant}}$ is already very well understood, and the probability of some structural properties that we discuss can be easily estimated for $\mathcal{P}_{n,m}^{\text{plant}}$ using standard probabilistic calculations. It then remains to find a reasonable "exchange rate" between $\mathcal{P}_{n,m}^{\text{plant}}$ and $\mathcal{P}_{n,m}^{\text{sat}}$. We use this approach to estimate the probability of "complicated" properties, which hold with extremely high probability in $\mathcal{P}_{n,m}^{\text{plant}}$. The other method is directly analyzing $\mathcal{P}_{n,m}^{\text{sat}}$, crucially overcoming the clause-dependency issue. This method tends to be more involved than the first one, and necessitates intricate counting arguments.

1.3 Paper's Structure

The rest of the paper is structured as follows. In Section 2 we discuss relevant structural properties that a typical formula in $\mathcal{P}_{n,m}^{\text{sat}}$ possesses. One consequence of this discussion will be a proof of Theorem 2. We then prove Theorem 1 in Section 3 by presenting an algorithm and showing that it meets the requirements of Theorem 1. Concluding remarks are given in Section 5. Due to lack of space most propositions are given without a proof which can be found in complete in the appendix. We do however include two complete proofs – one that uses the exchange-rate technique (Section 2.2), and one directly analyzing $\mathcal{P}_{n,m}^{\text{sat}}$ (Section 4).

2 Properties of a Random Instance from $\mathcal{P}_{n,m}^{\text{sat}}$

In this section we analyze the structure of a typical formula in $\mathcal{P}_{n,m}^{\text{sat}}$. One direct consequence of the discussion in this section is a proof of Theorem 2, another is that the algorithm that we describe in Section 3 meets the requirements of Theorem 1.

Here and throughout we think of m as $O(n \log n)$. Otherwise, typically a formula in $\mathcal{P}^{\text{sat}}_{n,m}$ with $m \geq C_0 n \log n$, C_0 some sufficiently large constant, has only one satisfying assignment (as implied by the proof of Proposition 6), and therefore by the definition of $\mathcal{P}^{\text{plant}}_{n,m}$, it holds that $\mathcal{P}^{\text{sat}}_{n,m}$ and $\mathcal{P}^{\text{plant}}_{n,m}$ are statistically close. Then a simple second moment calculation shows that the Majority Vote, discussed ahead, will reconstruct the satisfying assignment *whp* for $\mathcal{P}^{\text{plant}}_{n,m}$ in that regime. The interesting case remains $m = O(n \log n)$.

2.1 Setting the Exchange Rate

Let $\mathcal{A}$ be some property of CNF formulas (it would be convenient for the reader to think of $\mathcal{A}$ as a "bad" property). We start by determining the exchange rate for $Pr[\mathcal{A}]$ when moving from the planted distribution to the uniform one.

For a property $\mathcal{A}$ we use $Pr^{\text{uniform,m}}[\mathcal{A}]$ to denote the probability of $\mathcal{A}$ occurring under $\mathcal{P}^{\text{sat}}_{n,m}$, and $Pr^{\text{planted,m}}[\mathcal{A}]$ for $\mathcal{P}^{\text{plant}}_{n,m}$. The following lemma asserts the exchange rate $\mathcal{P}^{\text{plant}}_{n,m} \to \mathcal{P}^{\text{sat}}_{n,m}$. The proof is rather involved technically and embeds interesting results of their own – for example, bounding the expected number of satisfying assignments of a formula in $\mathcal{P}^{\text{sat}}_{n,m}$.

Lemma 1. *($\mathcal{P}^{\text{plant}}_{n,m} \to \mathcal{P}^{\text{sat}}_{n,m}$) Let $\mathcal{A}$ be some property of 3CNF formulas, then*

$$Pr^{\text{uniform,m}}[\mathcal{A}] \leq e^{ne^{-m/(3n)}} \cdot Pr^{\text{planted,m}}[\mathcal{A}].$$

Remark 1. Observe that the exchange rate between the planted distribution and the uniform is exponential in n. Thus we can use Lemma 1 whenever the "bad" event $\mathcal{A}$ happens with exponentially small probability in $\mathcal{P}^{\text{plant}}_{n,m}$. It is pretty straightforward to obtain an exchange rate of 2^n (which is far less useful, at least in our analysis); working out $e^{ne^{-m/(3n)}}$, though, demands more careful and non-trivial arguments. Full details are in the appendix.

2.2 The Majority Vote

For a 3CNF formula $\mathcal{F}$ and a variable x we let $N^+(x)$ be the set of clauses in $\mathcal{F}$ in which x appears positively (namely, as the literal x), and $N^-(x)$ be the set of clauses in which x appears negatively (that is, as $\bar{x}$). The Majority Vote assignment over $\mathcal{F}$, which we denote by MAJ, assigns every x according to the sign of $|N^+(x)| - |N^-(x)|$ (TRUE if the difference is positive and FALSE otherwise).

To show the usefulness of the Majority Vote in $\mathcal{P}^{\text{sat}}_{n,m}$ we work our way through $\mathcal{P}^{\text{plant}}_{n,m}$, and use the exchange-rate technique. Consider $\mathcal{F}$ in $\mathcal{P}^{\text{plant}}_{n,m}$, and let φ be its planted assignment. Consider a variable x whose assignment is w.l.o.g. $\varphi(x) = TRUE$. In every clause of $\mathcal{F}$ that contains x, x appears positively with probability 4/7, and negatively with probability 3/7. Therefore in expectation the sign of $|N^+(x)| - |N^-(x)|$ agrees with $\varphi(x)$. More formally, one can prove the following fact (see [21] for the complete proof):

Lemma 2. *Let $\mathcal{F}$ be distributed according to $\mathcal{P}^{\text{plant}}_{n,m}$ with $m \geq C_0 n$, C_0 a sufficiently large constant. Let F_{MAJ} be a random variable counting the number of variables in $\mathcal{F}$ on which MAJ disagrees with the planted assignment. There exists a constant $a_0 > 0$ (independent of m, n) and a positive monotonically increasing function f s.t. for every $a \geq a_0$ it holds that*

$$Pr[F_{MAJ} \geq e^{-m/(an)} n] \leq e^{-ne^{-m/(f(a)n)}}.$$

Proposition 1. *Let $\mathcal{F}$ be distributed according to $\mathcal{P}^{\text{sat}}_{n,m}$ with $m \geq C_0 n$, C_0 a sufficiently large constant. Then* whp *there exists a satisfying assignment φ of $\mathcal{F}$ that differs from MAJ on at most $e^{-\Theta(m/n)} n$ variables.*

Proof. Set $a_0 = f^{-1}(3)$ (f is the function promised in Lemma 2, $f^{-1}(3)$ is taken according to the denominator in $e^{ne^{-m/(3n)}}$ from Lemma 1), and $a_1 = 2a_0$. Let $\mathcal{A}$ be the following property: "there exists no satisfying assignment s.t. MAJ is at distance at most $e^{-m/(a_1 n)}n$ from it" (by distance we mean the Hamming distance). Using the exchange-rate technique we obtain:

$$Pr^{uniform,m}[\mathcal{A}] \underbrace{\leq}_{\text{Lemma 1}} e^{ne^{-m/(3n)}} \cdot Pr^{\text{planted,m}}[\mathcal{A}] \underbrace{\leq}_{\text{Lemma 2}} e^{ne^{-m/(3n)}} \cdot e^{-ne^{-m/(f(a_1)n)}}$$
$$= e^{n(e^{-m/(3n)} - e^{-m/(f(a_1)n)})} = o(1).$$

The last equality is by the choice of a_1 and the fact that f is increasing, that is $f(a_1) = f(2a_0) > f(a_0) = 3$ and therefore $e^{-m/(3n)} - e^{-m/(f(a_1)n)} < 0$.

2.3 The Discrepancy Property

A well known result in the theory of random graphs is that a random graph *whp* will not contain a small yet unexpectedly dense subgraph. This is also the case for $\mathcal{P}_{n,m}$ (when considering the graph induced by the formula). This property holds only with probability $1 - 1/poly(n)$ (for example, with probability $1/poly(n)$ a fixed clique on a constant number of vertices will appear). Thus the exchange-rate technique is of no use in this case (as the exchange rate factor is exponential in n). Overcoming the clause-dependency issue, using an intricate counting argument, we directly analyze $\mathcal{P}^{\text{sat}}_{n,m}$ to prove:

Proposition 2. *Let $\mathcal{F}$ be distributed according to $\mathcal{P}^{\text{sat}}_{n,m}$ with $m \geq C_0 n$, C_0 a sufficiently large constant. Then* whp *there exists* no *subset of variables U s.t.*

- $|U| \leq n/2000$*,*
- *There are $|U| \cdot \frac{m}{50n}$ clauses in $\mathcal{F}$ that contain two variables from U.*

The full proof is given in Section 4.

Remark 2. To see how Proposition 2 corresponds to the random graph context, consider the graph induced by the formula $\mathcal{F}$ (the vertices are the variables, and two variables share an edge if there exists some clause containing them both) and observe that every clause that contains at least two variables from U contributes an edge to the subgraph induced by U. Thus if we have many such clauses, this subgraph will be prohibitively dense. Since $\mathcal{F}$ is random so is its induced graph, and therefore the latter will typically not occur. In our case $\mathcal{F}$ is random but the clauses are dependent – making the analysis more complicated.

2.4 The Core Variables

We describe a subset of the variables, referred to as the *core variables*, which plays a crucial role in the understanding of $\mathcal{P}^{\text{sat}}_{n,m}$. Recall that a variable is said to be frozen in $\mathcal{F}$ if in every satisfying assignment it takes the same assignment. The notion of core captures this phenomenon. In addition, a core typically contains all but a small (though constant) fraction of the variables. This implies that a large fraction of the variables is frozen, a fact which must leave imprints on various structural properties of the formula. These imprints allow efficient heuristics to recover a satisfying assignment of the core. A second implication of this is an upper bound on the number of possible satisfying assignments, and on the distance between every such two. Thus the notion of core gives a catheterization of the cluster structure of the solution space (matching the properties described in Theorem 2).

Definition 1. *(support) Given a 3CNF formula $\mathcal{F}$, and some assignment ψ to the variables, we say that a literal x* supports *a clause C (in which it appears) w.r.t. ψ if x is the only literal that evaluates to true in C under ψ.*

Definition 2. *(core) A set of variables $\mathcal{H}$ is called a* **core** *of $\mathcal{F}$ w.r.t. to a satisfying assignment ψ, if the following three properties hold:*

- *Every variable $x \in \mathcal{H}$ supports at least $m/(5n)$ clauses in $\mathcal{F}[\mathcal{H}]$ w.r.t. ψ ($\mathcal{F}[\mathcal{H}]$ being the subformula containing the clauses where all three variables belong to $\mathcal{H}$).*
- *x appears in at most $m/(10n)$ clauses in $\mathcal{F} \setminus \mathcal{F}[\mathcal{H}]$.*

Remark 3. The proof of Theorem 2 (structure of the solution space) uses only the first property in Definition 2. However, since the core is also used for the algorithmic perspective, the second property is needed for the analysis of the algorithm.

Remark 4. The choice of $m/(5n)$ corresponds to slightly less than the expected support of a variable w.r.t. the planted assignment (which is roughly $3m/(14n)$), had the underlying probability space been $\mathcal{P}^{\text{plant}}_{n,m}$.

We proceed by asserting some relevant properties that such a core typically possesses.

Proposition 3. *Let $\mathcal{F}$ be distributed according to $\mathcal{P}^{\text{sat}}_{n,m}$ with $m \geq C_0 n$, C_0 a sufficiently large constant. Then* whp *there exits a satisfying assignment φ of $\mathcal{F}$ w.r.t. which there exists a core $\mathcal{H}$ and $|\mathcal{H}| \geq (1 - e^{-\Theta(m/n)})n$.*

The proof of Proposition 3 uses the exchange-rate technique, similar to the proof of Proposition 1 (a proof of Proposition 3 in the planted setting is given in [21], similar to Lemma 2). Details omitted.

The next proposition ties between the core variables and the property of the Majority Vote, and is crucial to the analysis of the algorithm. The proposition follows by noticing that φ in Lemma 2 and in its core-size counterpart is the same – the planted assignment. Thus, one can apply the exchange-rate technique on the combined property.

Proposition 4. *Let $\mathcal{F}$ be distributed according to $\mathcal{P}^{\text{sat}}_{n,m}$ with $m \geq C_0 n$, C_0 a sufficiently large constant. Then* whp *there exists a satisfying assignment φ s.t. the following two properties hold:*

- *MAJ differs from φ on at most $e^{-\Theta(m/n)}n$ variables*
- *There exists a core $\mathcal{H}$ w.r.t. φ as promised in Proposition 3.*

The next proposition characterizes the structure of the formula induced by the non-core variables. The connected components of a formula $\mathcal{F}$ are the sub-formulas $\mathcal{F}[C_1], \ldots, \mathcal{F}[C_k]$, where $C_1, C_2, \ldots, C_k$ are the the connected components in the graph induced by $\mathcal{F}$. Given a core $\mathcal{H}$ of $\mathcal{F}$ w.r.t. a satisfying assignment φ, we denote by $\mathcal{F}^{\varphi}_{out}(\mathcal{H})$ the subformula of $\mathcal{F}$ which is the outcome of the following procedure: set the variables $\mathcal{H}$ in $\mathcal{F}$ according to φ and simplify $\mathcal{F}$.

Proposition 5. *Let $\mathcal{F}$ be distributed according to $\mathcal{P}^{\text{sat}}_{n,m}$ with $m \geq C_0 n$, C_0 a sufficiently large constant. Let $\mathcal{H}$ be the core promised in Proposition 3. Then* whp *the largest connected component in $\mathcal{F}^{\varphi}_{out}(\mathcal{H})$ is of size $O(\log n)$.*

Proposition 5 also holds only with probability $1-1/poly(n)$, had $\mathcal{F}$ been distributed according to $\mathcal{P}^{\mathrm{plant}}_{n,m}$. Thus, similar to Proposition 2 the analysis is an involved counting argument (in this case even more complicated). Full details are in the appendix.

Lastly, we establish the "frozenness" property of the core variables. The proof uses Proposition 2 to show that there are no "close" satisfying assignments, and the exchange-rate technique to prove that there are no "far" ones. Full details are in the appendix.

Proposition 6. *Let $\mathcal{F}$ be distributed according to $\mathcal{P}^{\mathrm{sat}}_{n,m}$ with $m \geq C_0 n$, C_0 a sufficiently large constant. Let $\mathcal{H}$ be the core promised in Proposition 3. Then* whp *the assignment of $\mathcal{H}$ in all satisfying assignments of $\mathcal{F}$ is the same.*

3 Proof of Theorems 1 and 2

Theorem 2 is an immediate corollary of Propositions 3, 5 and 6 . Propositions 3 and 6 imply that all but a $e^{-\Theta(m/n)}n$ of the variables are frozen. Therefore, there are at most $2^{e^{-\Theta(m/n)}n}$ possible ways to set the assignment of the remaining variables. Furthermore, every two satisfying assignments of $\mathcal{F}$ can differ on the assignment of at most $e^{-\Theta(m/n)}n$ variables (that of the non-core variables). Proposition 5 completes the proof with the characterization of the formula induced by the non-frozen variables (which are a subset of the non-core ones).

Before proving **Theorem 1** we present an algorithm which we claim meets the requirements of Theorem 1. The algorithm is basically the one given in [14].

Remark 5. The versed reader in the area will notice some differences from the original algorithm in [14]. However, since we consider a different distribution than the one in [14], one can describe a simplified version of that algorithm (e.g., replace the spectral step with a Majority Vote).

SAT$(\mathcal{F})$
Step 1: Majority Vote
$\pi_1 \leftarrow$ Majority Vote over $\mathcal{F}$.
Step 2: Reassignment
for $i = 1$ to $\log n$
  for all $x \in V$
    if x supports less than $m/(5n)$ clauses w.r.t. π_i **then** $\pi_{i+1} \leftarrow \pi_i$ with x flipped.
  end for.
end for.
Step 3: Unassignment
set $\psi_1 = \pi_{\log n}$, $i = 1$.
while $\exists x$ s.t. x supports less than $m/(10n)$ clauses w.r.t. ψ_i
  set $\psi_{i+1} \leftarrow \psi_i$ with x unassigned.
  $i \leftarrow i + 1$.
end while.
Step 4: Exhaustive Search
Let ξ be the final partial assignment.
let A be the set of assigned variables in ξ.
exhaustively search $\mathcal{F}^{\xi}_{out}(A)$, component by component.

We now prove that the algorithm SAT meets the requirements of Theorem 1. We say that $\mathcal{F}$ is *typical* in $\mathcal{P}^{\mathrm{sat}}_{n,m}$ if Propositions 2, 4 and 5 hold for it. The discussion in Section

2 guarantees that *whp* $\mathcal{F}$ is typical. Therefore, to prove Theorem 1 it suffices to consider a typical $\mathcal{F}$ and prove that SAT (always) finds a satisfying assignment for $\mathcal{F}$.

We let $\mathcal{H}$ be the core promised in Proposition 3, and φ – the satisfying assignment w.r.t. which $\mathcal{H}$ is defined. In all the following propositions we assume $\mathcal{F}$ is typical (we don't explicitly state it every time for the sake of brevity). Similar propositions to Propositions 7–9 were proven in [14] for example. For completeness, all proofs are given in the appendix.

Proposition 7. *Let ψ_1 be the assignment defined in line 7 of SAT. Then ψ_1 agrees with φ on the assignment of all variables in $\mathcal{H}$.*

Proposition 8. *Let ξ be the partial assignment defined in line 12 of SAT. Then all assigned variables in ξ are assigned according to φ, and all the variables in $\mathcal{H}$ are assigned.*

Proposition 9. *The exhaustive search, Step 4 of SAT, completes in polynomial time with a satisfying assignment of $\mathcal{F}$.*

Theorem 1 then follows.

4 Proof of Proposition 2

Let V be the set of n variables, and let U be some fixed subset of V, $|U| = u$. Let H be a fixed formula over V with exactly $\frac{um}{50n}$ clauses s.t. each clause contains at least two variables from U. A formula $\mathcal{F}$ is said to be *H-poor* if it contains H as a sub-formula.

Furthermore, let $\mathcal{P}_H$ signify the set of all H-poor satisfiable formulas with exactly m clauses, and $\mathcal{A}$ the set of all satisfiable formulas with exactly m edges. Our first objective is to establish the following.

Lemma 3. $|\mathcal{P}_H| \le (em/n^3)^{um/(50n)}|\mathcal{A}|$.

This immediately implies that the probability of an H-poor formula in $\mathcal{P}^{\mathrm{sat}}_{n,m}$ is at most $(em/n^3)^{um/(50n)}$. Next take the union bound over all possible sub-formulas H (s.t. $|U| \le n/2000$ – as required by Proposition 2) to show that *whp* none is contained in a random $\mathcal{P}^{\mathrm{sat}}_{n,m}$ formula.

To prove Lemma 3 we shall set up an auxiliary bipartite graph $\mathcal{G}$ with vertex set $V(\mathcal{G}) = \mathcal{P}_H \cup \mathcal{A}$. This graph will have the property that the average degree of a vertex in $\mathcal{P}_H$ is Δ, while that of a vertex in $\mathcal{A}$ is Δ', where in addition $\Delta'/\Delta \le (em/n^3)^{um/(50n)}$. Since $\Delta \# \mathcal{P}_H = \Delta' \# \mathcal{A}$, by double counting, we thus obtain Lemma 3. We describe a nondeterministic procedure **P** that receives a formula $F \in \mathcal{P}_H$ and produces a new formula $F' \in \mathcal{A}$. In our auxiliary graph $\mathcal{G}$, we connect a right-side node F with a left-side one F', if F' can be obtained from F by applying **P** to F. **P** is the following procedure:

given a H-poor formula F do:
- Choose a set $\mathcal{C}$ of $um/(50n)$ fresh clauses (that are not yet in F)
- Obtain F' from F by removing all $um/(50n)$ clauses of H and adding $\mathcal{C}$.
- Output F' if it is satisfiable.

Therefore,
$$\Delta \ge \binom{n^3}{um/(50n)}.$$
This is because we have to choose $um/(50n)$ clauses out of at least $7\binom{n}{3} - m \ge n^3$ possible ones (since F was satisfiable to begin with, there are at least $7\binom{n}{3}$ clauses that are satisfied

by the assignment that satisfies F, and we can assume that $m = O(n \log n)$). Conversely, consider the following nondeterministic procedure to recover a formula F from F'. Out of m possible clauses in F', choose $um/(50n)$. Take them out, and reinstall the original clauses of H. Therefore,

$$\Delta' \leq \binom{m}{um/(50n)}.$$

Using standard bounds on the binomial coefficients, the required bound on Δ'/Δ is obtained and Lemma 3 follows.

We are now ready to bound the probability that a random formula $\mathcal{F}$ in $\mathcal{P}^{\text{sat}}_{n,m}$ violates the condition of Proposition 2. Using the union bound this probability is at most

$$\begin{aligned}
\sum_{u=1}^{n/2000} \binom{n}{u} \binom{8n\binom{u}{2}}{um/(50n)} \cdot \left(\frac{em}{n^3}\right)^{um/(50n)} &\leq \sum_{u=1}^{n/2000} \left(\frac{en}{u}\right)^u \left(\frac{600un^2}{m}\right)^{um/(50n)} \left(\frac{em}{n^3}\right)^{um/(50n)} \\
\leq \sum_{u=1}^{n/2000} \left(\frac{en}{u}\right)^u \left(\frac{1800u}{n}\right)^{uC_0/50} &\leq \sum_{u=1}^{n/2000} \left(\frac{en}{u} \cdot \frac{1800u}{n} \cdot \left(\frac{1800u}{n}\right)^{C_0/50-1}\right)^u \\
\leq \sum_{u=1}^{n/2000} \left(5400 \cdot \left(\frac{1800u}{n}\right)^{C_0/100}\right)^u &= o(1)
\end{aligned}$$

The last equality is due to $(u/n) \leq 1/2000$, so the last sum decreases faster than a geometric series with quotient and first element equal $1/poly(n)$, and therefore the whole sum is $o(1)$.

5 Discussion

Though $\mathcal{P}_{n,m}$ has a very simple description (fix $c, n > 0$ and choose $m = cn$ clauses uniformly at random out of $8\binom{n}{3}$ possible ones), and is very fundamental to understanding the hardness of 3SAT, it still baffles many researchers and altogether remains very poorly understood. In particular, the hardness of deciding if a random formula is satisfiable, and finding a satisfying assignment for a random formula, are both major open problems [10, 23].

Trying to shed some light on this problem we consider the uniform distribution over satisfiable 3CNF formulas, $\mathcal{P}^{\text{sat}}_{n,m}$, with clause-variable ratio greater than some sufficiently large constant. We characterize the typical structure of the solution space of such formulas, and show that a relatively simple efficient algorithm recovers *whp* a satisfying assignment of such formulas, thus asserting that almost all 3CNF formulas (when the clause-variable ratio is sufficiently large, yet possibly constant) are easy. To obtain our result we had to come up with new analytical tools that apply to a number of further NP-hard problems, including k-colorability. Our result also implies that the algorithmic techniques developed for random formulas from the planted distribution, e.g. [14, 12, 13, 21], can be extended to the significantly more natural uniform distribution.

Moreover, our result supports the assumption that the empirical hardness of some SAT distributions is mainly dictated by the structure of the solution space of a typical formula in that distribution. Specifically, the conjectured complicated clustering in the "hard" near-threshold regime versus the more "regular" structure that denser "easy" formulas possess.

References

1. D. Achlioptas and F. Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 130–139, 2006.
2. N. Alon and N. Kahale. A spectral technique for coloring random 3-colorable graphs. *SIAM J. on Comput.*, 26(6):1733–1748, 1997.
3. N. Alon, M. Krivelevich, and B. Sudakov. Finding a large hidden clique in a random graph. *Random Structures and Algorithms*, 13(3-4):457–466, 1998.
4. E. Ben-Sasson, Y. Bilu, and D. Gutfreund. Finding a randomly planted assignment in a random 3*CNF*. *manuscript*, 2002.
5. A. Blum and J. Spencer. Coloring random and semi-random k-colorable graphs. *J. of Algorithms*, 19(2):204–234, 1995.
6. A. Braunstein, M. Mezard, and R. Zecchina. Survey propagation: an algorithm for satisfiability. *Random Structures and Algorithms*, 27:201–226, 2005.
7. H. Chen. An algorithm for sat above the threshold. In *6th International Conference on Theory and Applications of Satisfiability Testing*, pages 14–24, 2003.
8. A. Coja-Oghlan, M. Krivelevich, and D. Vilenchik. Why almost all k-colorable graphs are easy. In *STACS*, 2007. to appear.
9. S. A. Cook. The complexity of theorem-proving procedures. In *Proc. 3rd ACM Symp. on Theory of Computing*, pages 151–158, 1971.
10. U. Feige. Relations between average case complexity and approximation complexity. In *Proc. 34th ACM Symp. on Theory of Computing*, pages 534–543, 2002.
11. U. Feige and R. Krauthgamer. Finding and certifying a large hidden clique in a semirandom graph. *Random Structures and Algorithms*, 16(2):195–208, 2000.
12. U. Feige, E. Mossel, and D. Vilenchik. Complete convergence of message passing algorithms for some satisfiability problems. In *Random*, pages 339–350, 2006.
13. U. Feige and D. Vilenchik. A local search algorithm for 3SAT. Technical report, The Weizmann Institute of Science, 2004.
14. A. Flaxman. A spectral technique for random satisfiable 3CNF formulas. In *Proc. 14th ACM-SIAM Symp. on Discrete Algorithms*, pages 357–363, 2003.
15. E. Friedgut. Sharp thresholds of graph properties, and the k-sat problem. *J. Amer. Math. Soc.*, 12(4):1017–1054, 1999.
16. R. G. Gallager. *Low-Density Parity-Check Codes.* SMIT Press, Cambridge, 1963.
17. J. Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
18. W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.
19. C. Hui and A. M. Frieze. Coloring bipartite hypergraphs. In *Proceedings of the 5th International IPCO Conference on Integer Programming and Combinatorial Optimization*, pages 345–358, 1996.
20. E. Koutsoupias and C. H. Papadimitriou. On the greedy algorithm for satisfiability. *Info. Process. Letters*, 43(1):53–55, 1992.
21. M. Krivelevich and D. Vilenchik. Solving random satisfiable $3cnf$ formulas in expected polynomial time. In *Proc. 17th ACM-SIAM Symp. on Discrete Algorithms*, pages 454–463, 2006.
22. L. Kučera. Expected behavior of graph coloring algorithms. In *Proc. Fundamentals of Computation Theory*, volume 56 of *Lecture Notes in Comput. Sci.*, pages 447–451. Springer, Berlin, 1977.
23. L. Levin. Average case complete problems. *SIAM J. Comput.*, 15(1):285–286, 1986.
24. M. Mezard, T. Mora, and R. Zecchina. Clustering of solutions in the random satisfiability problem. *Physical Review Letters*, 94:197–205, 2005.
25. T. Mora, M. Mezard, and R. Zecchina. Pairs of sat assignments and clustering in random boolean formulae, 2005.
26. J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference.* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1988.
27. B. Selman, H. A. Kautz, and B. Cohen. Local search strategies for satisfiability testing. In *Proceedings of the Second DIMACS Challange on Cliques, Coloring, and Satisfiability*, 1993.

A Proof of Lemma 1

Basically, what we show is that $Pr^{uniform,m}[\mathcal{A}]$ is at most the expected number of satisfying assignments that a formula in $\mathcal{P}^{\text{sat}}_{n,m}$ has times $Pr^{planted,m}[\mathcal{A}]$. To make this bound useful we need to bound this expectation (which is possibly an interesting result on its own right).

More formally, let $C_1(n)$ be the expected number of satisfying assignments of a formula in $\mathcal{P}^{\text{sat}}_{n,m}$, and $C_2(n)$ is defined similarly for $\mathcal{P}^{\text{plant}}_{n,m}$. Let t_i be the number of formulas on n variables and m clauses which have exactly i satisfying assignments. Let p_i be the probability that a formula with exactly i satisfying assignments is sampled from $\mathcal{P}^{\text{sat}}_{n,m}$, and let q_i be defined similarly for $\mathcal{P}^{\text{plant}}_{n,m}$. For a satisfying assignment φ, let $\Delta_{n,m,\varphi}$ be the number of formulas on n variables with m clauses that are satisfied by φ. Observe that due to symmetry $\Delta_{n,m,\varphi}$ is the same for every φ – thus we omit the φ subscript. In the above notation

$$p_i = \frac{t_i}{\sum_{j=1}^{2^n} t_j},$$

$$q_i = t_i \cdot \frac{i}{2^n} \cdot \frac{1}{\Delta_{n,m}}.$$

Further observe that

$$2^n \cdot \Delta_{n,m} = \sum_{i=1}^{2^n} i \cdot t_i.$$

This is because every formula with j satisfying assignments is counted exactly j times in the product $2^n \cdot \Delta_{n,m}$. Lastly,

$$C_1(n) = \sum_{i=1}^{2^n} i \cdot p_i = \frac{\sum_{i=1}^{2^n} i \cdot t_i}{\sum_{i=1}^{2^n} t_i},$$

$$C_2(n) = \sum_{i=1}^{2^n} i \cdot q_i = \frac{\sum_{i=1}^{2^n} i^2 \cdot t_i}{2^n \cdot \Delta_{n,m}} = \frac{\sum_{i=1}^{2^n} i^2 \cdot t_i}{\sum_{i=1}^{2^n} i \cdot t_i}.$$

Lemma 4. *Let $\mathcal{A}$ be some property of 3CNF formulas, then*

$$Pr^{\text{uniform,m}}[\mathcal{A}] \leq C_1(n) \cdot Pr^{\text{planted,m}}[\mathcal{A}].$$

Proof. First we obtain the following bound.

$$\frac{Pr^{\text{uniform,m}}[\mathcal{A}]}{Pr^{\text{planted,m}}[\mathcal{A}]} \leq \max_i \frac{p_i}{q_i}.$$

This follows from the following discussion. Let $T_{\mathcal{A}}$ be the set of satisfiable formulas for which property $\mathcal{A}$ holds.

$$Pr^{\text{uniform,m}}[\mathcal{A}] = \sum_{F \in T_{\mathcal{A}}} Pr^{\text{uniform,m}}[F], \;\; Pr^{\text{planted,m}}[\mathcal{A}] = \sum_{F \in T_{\mathcal{A}}} Pr^{\text{planted,m}}[F].$$

Now let $b = \max_i \frac{p_i}{q_i}$. For every satisfiable 3CNF F it holds that $Pr^{\text{uniform,m}}[F] \leq b \cdot Pr^{\text{planted,m}}[F]$. Therefore,

$$\sum_{F \in T_{\mathcal{A}}} Pr^{\text{uniform,m}}[F] \leq \sum_{F \in T_{\mathcal{A}}} b \cdot Pr^{\text{planted,m}}[F] = b \cdot \sum_{F \in T_{\mathcal{A}}} Pr^{\text{planted,m}}[F].$$

It now remains to estimate $\max_i \frac{p_i}{q_i}$.

$$\frac{Pr^{\text{uniform,m}}[\mathcal{A}]}{Pr^{\text{planted,m}}[\mathcal{A}]} \le \max_i \frac{p_i}{q_i} = \max_i \left(\frac{t_i}{\sum_{j=1}^{2^n} t_j} \right) \cdot \left(\frac{2^n \cdot \Delta_{n,m}}{i \cdot t_i} \right)$$
$$= \max_i \left(\frac{1}{i} \cdot \frac{\sum_{j=1}^{k^n} j \cdot t_j}{\sum_{j=1}^{2^n} t_j} \right) = \left(\frac{\sum_{j=1}^{2^n} j \cdot t_j}{\sum_{j=1}^{2^n} t_j} \right) \cdot \left(\max_i \frac{1}{i} \right) = C_1(n).$$

Since directly estimating $C_1(n)$ seems an intricate task, the following lemma is very useful.

Lemma 5. $C_1(n) \le C_2(n)$

Proof. To prove $C_1(n) \le C_2(n)$, one needs to prove that

$$\left(\sum_{i=1}^{2^n} i \cdot t_i \right)^2 \le \left(\sum_{i=1}^{2^n} t_i \right) \cdot \left(\sum_{i=1}^{2^n} i^2 \cdot t_i \right).$$

This is just Cauchy-Schwartz, $\left(\sum a_i \cdot b_i\right)^2 \le \left(\sum a_i^2\right) \cdot \left(\sum b_i^2\right)$, with $a_i = \sqrt{t_i}$ and $b_i = i \cdot \sqrt{t_i}$.

Corollary 1. *($\mathcal{P}^{\text{plant}}_{n,m} \to \mathcal{P}^{\text{sat}}_{n,m}$) Let $\mathcal{A}$ be some property of 3CNF formulas, then*

$$Pr^{\text{uniform,m}}[\mathcal{A}] \le C_2(n) \cdot Pr^{\text{planted,m}}[\mathcal{A}].$$

Proposition 10. *Let $\mathcal{F}$ be distributed according to $\mathcal{P}^{\text{plant}}_{n,m}$. Then $C_2(n) \le e^{ne^{-m/(3n)}}$ (for every ratio m/n).*

Proof. Let $\mathcal{F}$ be a planted instance with φ its planted assignment. Before estimating $C_2(n)$, we make the following observation. Let A_k be the event that an assignment at distance k from φ satisfies a clause which is also satisfied by φ. Easy counting arguments show that

$Pr[A_k] =$

$$\underbrace{1 \cdot \frac{\binom{n-k}{3}}{\binom{n}{3}}}_{\text{clause containing variables on which both assignments agree}} + \underbrace{\frac{6}{7} \cdot \left(1 - \frac{\binom{n-k}{3}}{\binom{n}{3}} \right)}_{\text{clause in which both disagree on at least one variable}} =$$

$$\frac{6}{7} + \frac{1}{7} \left(\frac{(n-k)(n-k-1)(n-k-2)}{n(n-1)(n-2)} \right) = \frac{6}{7} + \frac{1}{7} \left(1 - \frac{k}{n} \right)^3 + o(1) \le 1 - \frac{k}{3n}$$

$\mathcal{F}$ contains m clauses, thus the probability that an assignment ψ at distance k also satisfies $\mathcal{F}$ is $(Pr[A_k])^m = (1 - \frac{k}{3n})^m \le e^{-mk/(3n)}$.

Therefore,

$$C_2(n) \le \sum_{k=0}^{n} \binom{n}{k} \cdot 2^k \cdot e^{-mk/(3n)} \le \sum_{k=0}^{n} \binom{n}{k} \left(e^{-m/(3n)} \right)^k \cdot 1^{n-k}$$
$$= \left(1 + e^{-m/(3n)} \right)^n \le e^{ne^{-m/(3n)}}.$$

Remark 6. Observe that if for example $m/n \ge 4 \log n$, then $C_2(n) = 1 + o(1)$. That is, besides the planted assignment, one expects additional $o(1)$ satisfying assignments. Put differently, (using the Markov inequality) when $m/n \ge 4 \log n$ then *whp* there is only one satisfying assignment. This is then the regime where $\mathcal{P}^{\text{sat}}_{n,m}$ and $\mathcal{P}^{\text{plant}}_{n,m}$ are statistically close.

B Proof of Proposition 6

Let φ be the satisfying assignment w.r.t. which the core $\mathcal{H}$ is defined. To prove Proposition 6 we prove that there are no satisfying assignments "far" from φ (Proposition 11) and also no "close" ones (Proposition 13).

Proposition 11. *Let $\mathcal{F}$ be distributed according to $\mathcal{P}^{\text{sat}}_{n,m}$ with $m \geq C_0 n$, C_0 a sufficiently large constant, and let φ be the satisfying assignment w.r.t. which the core was defined. Then* whp *there exists* no *satisfying assignment of $\mathcal{F}$ that disagrees with φ on at least $n/2000$ variables in $\mathcal{H}$.*

To prove this proposition, we first prove that this property holds in $\mathcal{P}^{\text{plant}}_{n,m}$ with extremely high probability, and then use Lemma 1 to complete the proof. To prove that this property holds for $\mathcal{P}^{\text{plant}}_{n,m}$ we show that every assignment at distance $\geq n/2000$ from the planted assignment does not satisfy at least $m/5000$ clauses in $\mathcal{F}$, and since the core (which is defined w.r.t. the planted assignment in the planted case) typically contains all but $e^{-\Theta(m/n)}m$ of the clauses, there must be some clause in $\mathcal{F}[\mathcal{H}]$ which is left unsatisfied as well.

Proposition 12. *Let $\mathcal{F}$ be distributed according to $\mathcal{P}^{\text{plant}}_{n,m}$ with $m \geq C_0 n$, C_0 a sufficiently large constant, and let φ be its planted assignment. Then with probability at least $1 - 2^{-n}$, every assignment ψ at distance at least $n/2000$ from φ does not satisfy at least $m/10^5$ clauses in $\mathcal{F}$.*

Proof. The basic idea of the proof is to first calculate the expected number of clauses not satisfied by an assignment at distance at least $n/2000$ from φ, and show that this number is "much" higher than $m/10^5$, then show a concentration result.

Let ψ be an assignment at distance βn from φ. Let X_β be a random variable counting the number of clauses in $\mathcal{F}$ that ψ does not satisfy. Then we have:

$$E[X_\beta] \geq \beta m/14.$$

To see this observe that the probability that ψ satisfies a clause that is also satisfied by φ is at most:

$$\underbrace{1 \cdot \frac{\binom{n-\beta n}{3}}{\binom{n}{3}}}_{\text{clause containing variables on which both assignments agree}} + \underbrace{\frac{6}{7} \cdot \left(1 - \frac{\binom{n-\beta n}{3}}{\binom{n}{3}}\right)}_{\text{clause in which both disagree on at least one variable}}$$

$$= \frac{6}{7} + \frac{1}{7}\left(\frac{\binom{n-\beta n}{3}}{\binom{n}{3}}\right) \leq \frac{6}{7} + \frac{1}{7} \cdot e^{-3\beta n/n} \leq \frac{6}{7} + \frac{1}{7} \cdot \left(1 - \frac{\beta}{2}\right) = 1 - \beta/14$$

The second inequality uses the fact that $\beta \leq 1$, and the first inequality uses

$$\frac{\binom{a-x}{b}}{\binom{a}{b}} \leq \left(1 - \frac{b}{a}\right)^x \leq e^{-bx/a}.$$

Now set $\delta = 1 - \frac{14 \cdot 2000}{10^5} = 0.72$ (that is, $m/10^5 = (1-\delta) \cdot m/(2000 \cdot 14)$). Using for example the Chernoff bound (which is applicable since it is known that X_β is more concentrated than the corresponding quantity if the draws were made with replacement [18] – and then they would have been independent) one obtains that:

$$Pr[X_\beta \leq e^{-5b}m] \leq Pr[X_\beta \leq (1-\delta)\beta m/14] \leq Pr[X_\beta \leq (1-\delta)E[X_\beta]] \leq e^{-\delta^2 E[X_\beta]/3} \leq e^{-\beta m/6}.$$

Taking the union bound over all possible assignments, one obtains that the probability of an assignment at distance greater than $n/2000$ from φ not satisfying less than $m/10^5$ clauses is at most

$$\sum_{\beta=1/2000}^{1} \binom{n}{\beta n} 2^{\beta n} e^{-\beta m/6} \leq \sum_{\beta=1/2000}^{1} \left(\frac{2e}{\beta}\right)^{\beta n} e^{-\beta m/6} \leq \sum_{\beta=1/2000}^{1} \left(\frac{2e \cdot e^{-C_0/6}}{\beta}\right)^{\beta n} \leq$$
$$\sum_{\beta=1/2000}^{1} \left(6 \cdot e^{-C_0/6} \cdot 2000\right)^{\beta n} \leq \sum_{k=n/2000}^{n} \left(e^{-C_0/10}\right)^{k} \leq 2^{-n}.$$

The last inequality is due to the fact that the last sum is a geometric series with quotient $e^{-C_0/10}$, and the first element equals $e^{-C_0 n/20{,}000} = o(2^{-n})$ (for a sufficiently large C_0).

Proposition 13. *Let $\mathcal{F}$ be distributed according to $\mathcal{P}^{\mathrm{sat}}_{n,m}$ with $m \geq C_0 n$, C_0 a sufficiently large constant. Let $\mathcal{H}$ be some core of $\mathcal{F}$, and let φ be the underlying assignment. Then* whp *there exists* no *satisfying assignment of $\mathcal{F}$ that disagrees with φ on at most $n/2000$ variables of $\mathcal{H}$.*

Proof. Assume that Proposition 2 (no small yet dense sub-formulas) indeed holds for $\mathcal{F}$, which is the case *whp*. Fix an arbitrary t-core $\mathcal{H}$, $t = m/(5n)$, and let ψ be a "bad" satisfying assignment of $\mathcal{F}$ – that is, ψ disagrees with φ on the assignment of at most $n/2000$ core variables (φ is the satisfying assignment w.r.t. which $\mathcal{H}$ is defined). Let x be some variable on which they disagree (if none exists, then we are done). Now consider all the clauses that x supports w.r.t. φ where all variables belong to $\mathcal{H}$. It must be that every such clause contains another *core* variable on which ψ and φ disagree (since every such clause is satisfied by φ, and the literal of x is false w.r.t. φ). Put differently, let U be the set of core variables on which ψ and φ disagree. By the discussion above, there are $|U| \cdot \frac{m}{5n}$ clauses each containing two variables from U (no clause was counted twice since the supporter of a clause is unique by definition). By the contradiction assumption $|U| \leq n/2000$, this however contradicts Proposition 2.

C Proof of Proposition 7

Let B_i be the set of core variables whose assignment in ψ_i disagrees with φ at the beginning of the i^{th} iteration of the main for-loop – line 2 in SAT. It suffices to prove that $|B_{i+1}| \leq |B_i|/2$ (if this is true, then after $\log n$ iterations $B_{\log n} = \emptyset$). Observe that by Proposition 4 $|B_0| \leq n/2000$. By contradiction, assume that not in very iteration $|B_{i+1}| \leq |B_i|/2$, and let j be the first iteration violating the inequality – $|B_{j+1}| \geq |B_j|/2$. Consider a variable $x \in B_{j+1}$. If also $x \in B_j$, this means that x's assignment was not flipped in the j^{th} iteration, and therefore, x supports at least $m/(5n)$ clauses w.r.t. ψ_j. By the second item in the definition of a core, at least $m/(5n) - m/(10n) = m/(10n)$ of these clauses contain only core variables. Since the literal of x is true in all these clauses, but in fact should be false under φ, each such clause must contain another variable on which φ and ψ_j disagree, that is another variable from B_j. If $x \notin B_j$, this means that x's assignment was flipped in the j^{th} iteration. This is because x supports less than $m/(10n)$ clauses w.r.t. ψ_j. Since x supports at least $m/(5n)$ clauses w.r.t. φ, it must be that in at least $m/(5n) - m/(10n) = m/(10n)$ of them, the literal of some other core variable evaluates to true (rather than false, as it should be w.r.t. φ). For conclusion, let $U = B_j \cup B_{j+1}$. Then there are at least $m/(10n) \cdot |B_{j+1}|$ clauses containing at least two variables from U. Now if $|B_{j+1}| \geq |B_j|/2$, then $m/(10n) \cdot |B_{j+1}| \geq m/(20n) \cdot |U|$, contradicting Proposition 2.

D Proof of Proposition 8

By the definition of core – the Majority Vote assignment sets the core variables correctly in ψ_1 (the assignment defined in line 2 in the algorithm SAT) – that is according to φ (φ is the satisfying assignment w.r.t. which $\mathcal{H}$ is defined). Furthermore, by the definition of core, every core variable supports at least $m/(5n)$ clauses w.r.t. φ, and also w.r.t. ψ_1 (the assignment at hand before the unassignment step begins). Hence all core variables survive the first round of unassignment. By induction it follows that the core variables survive all rounds. Now suppose by contradiction that not all assigned variables are assigned according to φ when the unassignment step ends. Let U be the set of variables that remain assigned when the unassignment step ends, and whose assignment disagrees with φ. Every $x \in U$ supports at least $m/(10n)$ clauses w.r.t. to ξ (the partial assignment defined in line 7 of SAT), but each such clause must contain another variable on which ψ and φ disagree (since the clause is satisfied by φ, and φ falsifies the literal of x in each such clause). Thus, we have $|U| \cdot \frac{m}{10n}$ clauses each containing at least two variables from U. Since $U \cap \mathcal{H} = \emptyset$ (by the first part of this argument) and $|\mathcal{H}| \geq (1 - e^{-\Theta(m/n)})n$ (by Proposition 4) it follows that $|U| \leq e^{-\Theta(m/n)}n < n/2000$, contradicting Proposition 2.

E Proof of Proposition 9

By Proposition 8, the partial assignment at the beginning of the exhaustive search step is patrial to some satisfying assignment of the entire formula. Therefore the exhaustive search will succeed. Further observe that the unassigned variables are a subset of the non-core variables (Proposition 8). Proposition 5 then guarantees that the running time of the exhaustive search will be at most polynomial.

F Proof of Proposition 5

Let $d = \frac{m}{n}$. Let us say that a 3CNF F is *bounded* if the following conditions hold.

B1. For all $X \subset V$ such that $\#X \leq n/d^2$ there are at most $10\#X$ clauses containing more than two variables from X.
B2. Every variable appears in at most $\ln^2 n$ clauses.
B3. Let H be a subformula of $\mathcal{F}$ on $\#V(H) \geq (1 - d^{-10})n$ variables, so that every variable in H supports at least $d/50$ clauses in H. Then H is uniquely satisfiable.

Moreover, we call F *ε-feasible* if F has an induced subformula H with the following properties.

F1. $\#V(H) \geq (1 - \varepsilon \exp(-\sqrt{d}))n$ and $\#H$ contains at least $(1 - \varepsilon)m$ clauses.
F2. There exists a satisfying assignment φ of $\mathcal{F}$ so that every variable $x \in H$ supports at least $(1 - \varepsilon)d/30$ clauses in H w.r.t. φ.
F3. Every variable in H appears in at most εd clauses where not all variables belong to H.
F4. H is uniquely satisfiable.

If H, K are two induced subformulas of F that satisfy F1–F4, then the same is true for $H \cup K$. Therefore, F has a unique maximal induced subformula that enjoys F1–F4; this subformula will be denoted by F_ε in the sequel. Also observe that if F is ε-feasible then it is also ε'-feasible for $\varepsilon' \leq \varepsilon$ and it holds that $V(F_\varepsilon) \subseteq V(F_{\varepsilon'})$ (by $V(F)$ we denote the variable set of a formula F).

Lemma 6. *For any fixed $\varepsilon > 0$ $\mathcal{P}^{\mathrm{sat}}_{n,m}$ is bounded and ε-feasible* whp.

The Lemma follows from the discussion in Section 2. Also observe that F_ε, with, say, $\varepsilon \leq 0.02$, is a core of F according to Definition 2. Therefore when reading Proposition 5, one can think of the core $\mathcal{H}$ as $\mathcal{F}_{0.02}$.

Let $T \subset V$ be a set of size $t = \lceil \log n \rceil$, and let τ be a tree with vertex set T. Let F_τ be a fixed collection of clauses such that each edge of τ is induced by some clause of F_τ. We say that a clause set F_τ is *minimal* w.r.t. τ if by deleting a clause from F_τ, τ is not induced by F_τ anymore. By the definition of minimality, $|F_\tau| \leq |E(\tau)| = |V(\tau)| - 1$ (as τ is a tree). Moreover, let us call $\mathcal{F}$ (T, τ, F_τ)*-poor* if

- $\mathcal{F}$ is bounded,
- $\mathcal{F}$ is 0.01-feasible,
- $\mathcal{F}$ contains F_τ as a subformula,
- $V(\tau)$ does not intersect $\mathcal{F}_{0.02}$.

Denote by $\mathcal{G}$ the set of all satisfiable 3CNF formulas with variable set $V = \{x_1, \ldots, x_n\}$ and exactly m clauses, and let $\mathcal{P}(T, \tau, F_\tau)$ signify the set of all (T, τ, F_τ)-poor formulas $F \in \mathcal{G}$. Below we shall establish the following.

Lemma 7. *We have* $\#\mathcal{P}(T, \tau, F_\tau) \leq \left(e^{-\Theta(tm/n)} \cdot \left(\frac{m}{n^3}\right)^{\#F_\tau}\right) \#\mathcal{G}$.

Where by $\#F_\tau$ we denote the number of clauses in F_τ. Before we prove Lemma 7, let us note that it implies Proposition 5 immediately (thinking of the core $\mathcal{H}$ as $\mathcal{F}_{0.02}$). First let us establish the following two facts.

Theorem 3. *(Cayley) The number of spanning trees of K_r is r^{r-2}*

Lemma 8. *Let τ be a fixed tree on t vertices. The number of minimal sets of clauses that induce τ is at most*

$$\sum_{s=0}^{t/2} \binom{t}{2} 7^{t-s-1} n^{t-2s-1}$$

Proof. Let F_τ be a minimal spanning set. We proved that a minimal set contains at most $t-1$ clauses. Every clause may cover up to 2 edges (otherwise τ contains a cycle). We can extend a tree to a minimal set by having s clauses cover 2 edges of τ and $\tau - 1 - 2s$ clauses cover one edge. The total number of clauses in F_τ is $t - s - 1$. For every clause there are $(2^3 - 1)$ ways to set the signs of the variables. For every clause that covers one edge, there are $(n-2)$ ways to choose the third variable. A clause that covers 2 edges is determined uniquely by the two edges. We have at most $\binom{t}{2}$ ways to form the clauses that cover 2 edges. s may range from 0 to $t/2$.

Proof of Proposition 5. Now observe that the probability of some tree τ of size $\log n$ not intersecting $\mathcal{F}_{0.02}$ is at most

$$\sum_{s=0}^{\log n/2} \underbrace{\binom{n}{\log n} \cdot (\log n)^{\log n - 2}}_{\text{number of trees of size } \log n} \cdot \underbrace{7^{\log n - s - 1} \cdot n^{\log n - 2s - 1} \binom{\log n}{2}}_{\text{extend to a minimal set}} \cdot \underbrace{\left(\frac{m}{n^3}\right)^{\log n - s - 1} \cdot n^{-\Theta(m/n)}}_{\text{tree outside } \mathcal{F}_{0.02} \text{ (Lem. 7)}} \leq$$

$$\sum_{s=0}^{\log n/2} \left(\frac{e \cdot n}{\log n}\right)^{\log n} \cdot (\log n)^{\log n} \cdot n^{\log 7 + \log m/n} \cdot n^{\log n - 2s - 1 - 2(\log n - s - 1)} \cdot n^{-\Theta(m/n)} =$$

$$\sum_{s=0}^{\log n/2} n^{-\Theta(m/n) + \log 7 + \log m/n + 1} \cdot n^{\log n} \cdot n^{-\log n + 1} = \sum_{s=0}^{\log n/2} n^{-\Theta(m/n) + log 7 + \log m/n + 2}$$

$$= n^{-\Theta(m/n)}$$

Thus, the remaining task is to prove Lemma 7. To this end we fix a set T of variables, a tree τ on the variables T, and a minimal inducing set of clauses F_τ. We set up a bipartite auxiliary graph $\mathcal{A} = \mathcal{A}(T, F_\tau, \tau)$ with vertex set $V(\mathcal{A}) = \mathcal{P}(T, F_\tau, \tau) \oplus \mathcal{G}$; for brevity we set $\mathcal{P} = \mathcal{P}(T, F_\tau, \tau)$. The auxiliary graph will enjoy the following property. In $\mathcal{A}$ every vertex $F \in \mathcal{P}$ has degree at least Δ, while every vertex $F' \in \mathcal{G}$ has degree at most Δ' s.t.

$$\Delta' \leq \left(e^{-\Theta(tm/n)} \cdot \left(\frac{m}{n}\right)^{\#F_\tau}\right) \Delta \tag{1}$$

Since $\Delta \#\mathcal{P}(T, \tau) \leq \#E(\mathcal{A}) \leq \Delta' \#\mathcal{G}$, Lemma 7 follows directly from (1).

To describe the construction of $\mathcal{A}$ we let I be the set of all $x \in T$ that appear in at most 6 clauses in F_τ; then $\#I \geq t/2$, because $\#F_\tau \leq \#V(\tau) - 1$ (as τ is a tree and F_τ is minimal). Let φ be the unique satisfying assignment of $F_{0.02}$. We define the following partition of I:

$I_1(F) = \{x \in I : x$ appears in at least $0.02d$ clauses where all variables belong to $V \setminus F_{0.02}\}$,

$I_2(F) = \{x \in I : x$ supports at most $(1 - 0.02)d/30$ clauses w.r.t. φ where the other two variables belong to $F_{0.02}\} \setminus I_1(F)$.

If F is (T, F_τ, τ)-poor, then all variables $x \in I$ are outside of the 0.02-core $F_{0.02}$; hence, due to F1–F4 we have $I = I_1(F) \cup I_2(F)$. Thus, we decompose $\mathcal{P}$ into two parts $\mathcal{P}_1 = \{F \in \mathcal{P} : \#I_1(F) \geq 0.15t\}$, $\mathcal{P}_2 = \mathcal{P} \setminus \mathcal{P}_1$ (that is, $\#I_2(F)$ is at least $0.85t$).

As a next step, we will construct two subgraphs $\mathcal{A}_1$, $\mathcal{A}_2$ of $\mathcal{A}$, both of which consist of the $\mathcal{P}_i$-$\mathcal{G}$-edges of $\mathcal{A}$. Thus, $\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2$, so that (1) will be a consequence of the following statement. In $\mathcal{A}_j$ every vertex $F \in \mathcal{P}_j$ has degree at least Δ_j, while every vertex $F' \in \mathcal{G}$ has degree at most Δ'_j where

$$\Delta'_j \leq \left(e^{-\Theta(tm/n)} \cdot \left(\frac{m}{n^3}\right)^{\#F_\tau}\right) \Delta_j \quad (j = 1, 2). \tag{2}$$

In the remainder of this section we present the constructions of $\mathcal{A}_1, \mathcal{A}_2$ and establish (2). To facilitate these constructions we say that a triplet $\{x, y, z\}$ of variables is *compatible* if there is no clause in F involving all three variables, and x, y, z lie in $F_{0.01}$. Moreover, we say that a set F of triplets of variables is compatible if every triplet in F is compatible and no variable x occurs in more than one triplet.

Lemma 9. *Let $F \in \mathcal{P}$ and let $1 \leq s \leq n^{0.1}$. There exist $\binom{n^3/4}{s}$ compatible sets F of size s.*

Proof. Let φ signify the unique satisfying assignment of $F_{0.01}$, and let $\mathcal{C}$ be all the clauses over the variables of $F_{0.01}$ that are satisfied by φ. Since F satisfies F1, $\mathcal{C}$ has at least $7\binom{0.99n}{3} \geq n^3/2$

clauses. Furthermore, let S be a set of s clauses of $\mathcal{C}$ chosen uniformly at random. Then the probability that S does not contain a clause of F is

$$\binom{\#\mathcal{C}-m}{s}\binom{\#\mathcal{C}}{s}^{-1} = \prod_{j=0}^{s-1} 1 - \frac{m}{\#\mathcal{C}-j} = 1-o(1),$$

because $\#\mathcal{C} = \Omega(n^3)$, while $ms = o(n^3)$. Moreover, the probability that a specific variable x occurs twice in S is at most

$$\binom{7n^2}{2}\binom{\#\mathcal{C}}{s-2}\binom{\#\mathcal{C}}{s}^{-1} \leq O(s^2n^{-2}) = o(n^{-1}).$$

Hence, by the union bound with probability $1-o(1)$ a randomly chosen S will touch no variable x more than once. Thus, with probability $1-o(1)$ a randomly chosen S is compatible, so that the number of compatible sets is $\geq (1-o(1))\binom{\#\mathcal{C}}{s} \geq \binom{n^3/4}{s}$.

Construction of $\mathcal{A}_1$***.*** The construction of $\mathcal{A}_1$ is based on the following observation.

Lemma 10. *Suppose that* $F \in \mathcal{P}_1$*. There exist sets* $U \subset I_1(F)$*,* $\#U = \lceil 0.1t \rceil$*, and* $W \subset V \setminus (\tau \cup F_{0.02})$ *such that for every* $x \in U$*,* x *appears in at least* $d/1000$ *clauses where the other two variables belong to* W*, and for every* $y \in W$*,* y *appears in at most* 1000 *clauses where the other two variables belong to* U*.*

Proof. Let $J \subset I_1(G)$ be a set of size $0.15t$, and let K be the set of all variables $w \in V \setminus (F_{0.02} \cup \tau)$ s.t. there exists a clause in F containing w and some variable from J. Moreover, let $L \subset K$ be the set of all $w \in K$ such that the number of clauses containing w and at least one more variables in J is at least 1000. Then the boundedness property of F implies that $\#L \leq 0.01t$. Furthermore, letting Q be all the variables $w \in J$ such that the number of clauses containing w and at least one more variable from L is at least 1000. Then we have $\#Q \leq 0.01t$ (once more due to the boundedness of F). Now, let $U = J \setminus Q$ and $W = K \setminus L$. Then each $w \in W$ appears in at most 1000 clauses where the other two variables are in U. Moreover, if $v \in U$, then the number of clauses that contain v and two variables in K is at least $0.02d - 6 - 10^4 \geq 0.015d$ Furthermore, U has the required size.

Our objective is to associate with each $F \in \mathcal{P}_1$ a large number of "target graphs" $F' \in \mathcal{G}$ such that no F' occurs as a target graph too frequently. To this end, we consider the following nondeterministic procedure that maps F to a target graph F'. For each possible outcome F' we include the edge $\{F, F'\}$ into $\mathcal{A}_1$.
Set $\gamma = \lceil d/1000 \rceil$ and $u = \lceil 0.1t \rceil$.

C1. Choose a compatible set C of size $\#F_\tau + \gamma u$.
C2. Choose sets U and W as in Lemma 10.
C3. For each $x \in U$ choose a set N_x of γ clauses of the form $(\ell_x \vee \ell_y \vee \ell_z)$, $y, z \in W$.
C4. Obtain F' from F by removing the clauses of F_τ along with the clauses of C3 and adding the clauses from C.

Lemmas 9 entails that the number of formulas F' that can be obtained from each F via the above procedure is at least

$$\Delta_1 = \binom{n^3/4}{\#F_\tau + \gamma u} \tag{3}$$

(because there are at least this many choices in step C1, and one can always set the signs of variables in each new clause so that the obtained formula remains satisfiable – as we started with a satisfiable one). Conversely, to recover F from F', we consider the following nondeterministic procedure.

R1. Choose a set C' of $\#F_\tau + \gamma u$ clauses of F'.
R2. Choose a set $U' \subset T$ of size u.
R3. For each such $x \in U'$ choose a set N'_x of γ pairs of variables outside of $F_{0.015}$, and a way to set the signs of the variables in each clause.
R4. Output the formula F'' obtained from F' by removing the clauses C' and adding the ones $(\ell_x \vee \ell_y \vee \ell_z)$, $x \in U'$, $y, z \in N'_x$ along with the clauses of F_τ.

Lemma 11. *If $\{F, F'\}$ is an edge of $\mathcal{A}_1$, then F' is 0.015-feasible and the process R1–R4 applied to F' can yield the output $F'' = F$.*

Proof. Let C, U, W, and $(N_x)_{x\in U}$ be the sets chosen by C1–C4 to obtain F' from F. If R1–R4 chooses $F' = F$, $U' = U$, $N'_x = N_x$ for all $x \in U$, then the outcome will be $F'' = F$. Thus, we just need to show that it is feasible for R1–R4 to choose $N'_x = N_x$, i.e., that F' is 0.015-feasible and the vertices in N'_x do not belong to $F'_{0.015}$.

It suffices to show that $V(F'_{0.015}) \subseteq V(F_{0.02})$ (since all the variables in N_x lie outside $F_{0.02}$, and in particular outside $F'_{0.015}$).

To see that F' is 0.015-feasible, let X be the variable set of $F_{0.01}$. We claim that X satisfies F1–F4 with respect to F' with $\varepsilon = 0.01$. F1 is an immediate consequence of the fact that F is 0.01-feasible. Moreover, as C4 adds a compatible set C and only removes clauses that contain variables outside of X, the unique satisfying assignment of $F_{0.01}$ remains then unique for the set X in F', whence F2–F4 follow. Thus, F' is indeed 0.01-feasible, and hence 0.015-feasible as well.

Finally, to show that the variable set Y of $F'_{0.015}$ is contained in that of $F_{0.02}$, we show that Y is 0.02-feasible in F. Requirement F_1 is satisfied since F' is 0.015-feasible (and $0.015 > 0.02$). Further observe that if $F'[Y]$ is uniquely satisfiable then so is $F[Y]$, this is because when moving from F' to F one can either add clauses, or remove clauses that are uniquely satisfied in F to begin with (therefore removing them incurs no addition of satisfying assignments) – thus requirement F4 follows. Now consider the $(1 - 0.015)d/30$ clauses that every $y \in Y$ supports in $F'[Y]$, then since $F' \setminus F$ contains at most 1 clause involving y (as C is a compatible set), it holds that y supports at least $(1-0.015)d/30 - 1 \geq (1-0.02)d/30$ clauses w.r.t. $F[Y]$ (requirement F_2). Lastly, observe that if y appears in at most $0.015d$ clauses in F' where not all variables belong to Y, then in F this number could have been at most $0.015d + \gamma \leq 0.02d$, by the choice of γ (this establishes requirement F_3).

Lemma 12. *If F' is an outcome of C1–C4 for some $F \in \mathcal{P}_1$, then the number of possible nondeterministic choices in the R1–R4 is at most $\Delta'_1 = 2^{\#T}\binom{m}{\#F_\tau+\gamma u}\binom{\exp(-\sqrt{d})n^2}{\gamma}^u$.*

Proof. The first factor accounts for the number of ways to choose F'. Moreover, there are clearly at most $2^{\#T}$ ways to choose U' (recall that in our setting $\#T = \log n$). To bound the number of choices of R3, note that for each $x \in U'$ there are at most $\binom{(n-|V(F'_{0.015})|)^2}{\gamma}$ ways to choose the set N'_x. By the definition of $F'_{0.015}$ (requirement F_1) it holds that $|V(F'_{0.015})| \geq n(1 - \exp(-\sqrt{d}))$.

Finally, combining (3) with Lemma 11 and 12, and using standard estimates for the binomial coefficients, one obtains (2) for $j = 1$.

Construction of $\mathcal{A}_2$. As in the construction of $\mathcal{A}_1$ we consider a nondeterministic procedure that maps $F \in \mathcal{P}_2$ to $F' \in \mathcal{G}$. Let $u = \lceil 0.1t \rceil$ and $\gamma = \lceil 10^{-9}d \rceil$. Let φ be the unique satisfying assignment of $F_{0.01}$.

C1. Choose a compatible set C of size $\#F_\tau$.

C2. Choose a subset $U \subset I_2(F)$ of size u.

C3. Choose a set of clauses $M \subset F_{0.01}$ of size γu s.t. every two clauses in M are variable-disjoint, and no variable x appears in more than 100 clauses with a variable from M. Moreover, for each $x \in U$ choose a set N_x of clauses $C = (\ell_x \vee \ell_y \vee \ell_z)$ s.t. $y, z \in F_{0.02}$, there in no clause in F that contains both x, z or both x, y, and no variable in N_x occurs in M. Set the negation in every such C so that x supports C w.r.t. φ. Moreover, the sets $(N_x)_{x \in U}$ should be pairwise disjoint.

C4. Obtain F' from F by removing the clauses of F_τ and the set M, adding the clauses of C, and adding all the clauses $(N_x)_{x \in U}$.

For each $F \in \mathcal{P}_2$ and each possible outcome F' of C1–C4 we include the edges $\{F, F'\}$ into $\mathcal{A}_2$. The following lemma provides a lower bound on the degree of $F \in \mathcal{P}_2$ in $\mathcal{A}_2$.

Lemma 13. *Each $F \in \mathcal{P}_2$ has at least $\Delta_2 = \frac{1}{2}\binom{n^3/4}{\#F_\tau} \cdot \binom{m/2}{\gamma u} \cdot \binom{n^2\left(1-e^{-\sqrt{d}}\right)}{\gamma}^u$ images F'.*

Proof. By Lemma 9 there are $\binom{n^3/4}{\#F_\tau}$ ways to choose C. Furthermore, property F1 implies that $F_{0.01}$ contains at least $m/2$ clauses. Moreover, since every variable appears in at most $\ln^2 n$ clauses in F, property B2, and the boundedness of F (property B1 – which assures that for every M of size γu there are not too many variables that appear in at least 100 clauses with some other variable from M) then $F_{0.01}$ has at least $(1 - o(1))\binom{m/2}{\gamma u}$ sets M of size γu. Finally, since $\#V(F_{0.02}) \geq (1 - 0.02e^{-\sqrt{d}})n$ by F2, there are at least $\binom{n^2\left(1-e^{-\sqrt{d}}\right)}{\gamma}^u$ ways to choose the sets $(N_x)_{v \in U}$.

Conversely, we consider the following nondeterministic procedure for obtaining a formula F'' from an outcome F' of C1–C4.

R1. Choose a set of clauses C' from F' of size $\#F_\tau$.

R2. Determine the unique satisfying assignment φ of $F'_{0.015}$. Then, choose a set $U' \subset T$ of size u. Moreover, choose a set M' of γu clauses out of all possible ones.

R3. For each $x \in U'$ out of the clauses that x supports in $F'_{0.015}$, choose γ such clauses and set N'_x to be that set of these clauses.

R4. Obtain a formula F'' from F' by removing C' and all clauses in N'_x, and adding the clauses of F_τ and M'.

Lemma 14. *If $\{F, F'\}$ is an edge of $\mathcal{A}_2$, then F' is 0.015-feasible and the process R1–R4 applied to F' can yield the output $F'' = F$.*

Proof. Suppose that F' has been obtained from F by choosing the sets M, U, the sets $(N_x)_{x \in U}$, and the compatible set C. To recover $F'' = F$, we shall prove that F' is 0.015 feasible and that the process R1–R4 can choose $M' = M$, $C' = C$, and $N'_x = N_x$.

To see that F' is 0.015-feasible, let X be the variable set of $F_{0.01}$. We claim that X satisfies F1–F4 with respect to F' with $\varepsilon = 0.015$. For F1 is an immediate consequence of the fact that F is 0.01-feasible. Moreover, as C1 adds a compatible set C, and the N_x's are variable disjoint, and of size at most γ, then every $x \in X$ appears in at most $0.01d + \gamma + 1 \leq 0.015d$ clauses where not all variables belong to X. Every $x \in X$ supports at least $(1 - 0.01)d/30$ clauses where all other variables belong to X, and since M removes at most 100 clauses per variable, and C is compatible, then x supports at least $(1-0.01)d/30-100-1 \geq (1-0.015)d/30$ clauses (w.r.t. the same satisfying assignment, which remains satisfying for F' as well). Lastly, $F'[X]$ is uniquely satisfiable by property B3. This in turn implies that the same unique assignment satisfies $F[X]$ (otherwise, if φ' is the different assignment that satisfies $F[X]$ then φ' also satisfies $F'[X]$, as in $F'[X]$ we either removed clauses or added clauses which are satisfied by φ' – contradicting the uniqueness).

It is clear that R1–R4 can choose $C' = C$ and $M' = M$ since we had no restrictions in step R2. It only remains to show that it is feasible for the reconstruction procedure to have chosen $N'_x = N_x$. To this end, it suffices to show that $V(F'_{0.015}) \subseteq V(F_{0.02})$ (since all the variables in N_x lie outside $F_{0.02}$, and then in particular outside $F'_{0.015}$, and the set of clauses that x supports in $F'_{0.015}$ will then contain the set of clauses that it supports in $F_{0.02}$ – as the unique satisfying assignment is the same). Let Y be the set of variables of $F'_{0.015}$. We show that Y is 0.02-feasible in F. Requirement F_1 is satisfied since F' is 0.015-feasible (and $0.015 > 0.02$). Further observe that if $F'[Y]$ is uniquely satisfiable then so is $F[Y]$, this is because when moving from $F'[Y]$ to $F[Y]$ one can either add clauses, or remove clauses that are uniquely satisfied in F to begin with (therefore removing them incurs no addition of satisfying assignments) – thus requirement F4 follows. Now consider the $(1 - 0.015)d/30$ clauses that every $y \in Y$ supports in $F'[Y]$, then since $F' \setminus F$ contains at most 1 clause involving y (as C is a compatible set), it holds that y supports at least $(1-0.015)d/30-1 \geq (1-0.02)d/30$ clauses w.r.t. $F[Y]$ (requirement F_2). Lastly, observe that if y appears in at most $0.015d$ clauses in F' where not all variables belong to Y, then since the N_x are of size γ and variable-disjoint – in F this number could have been at most $0.015d + \gamma \leq 0.02d$ (by the choice of γ) – this establishes requirement F_3.

In the light of Lemma 14 we can bound the degrees of $F' \in \mathcal{G}$ in $\mathcal{A}_2$ as follows.

Lemma 15. *If G' has been obtained from G via C1–C4, then during R1–R4 there are at most $\Delta'_2 = \binom{m}{\#F_\tau} 2^t \binom{8\binom{n}{3}}{\gamma u} \binom{d/30}{\gamma}^u$ ways to choose F', the sets N'_x, and M'.*

Proof. There are exactly $\binom{m}{\#F_\tau}$ ways to choose C' and at most 2^t ways of choosing U'. Furthermore, there are at most $\binom{8\binom{n}{3}}{\gamma u}$ ways to choose M'. Finally, since each $x \in U'$ has at most $(1 - 0.02)d/30 + \gamma u \leq d/30$ clauses that it supports w.r.t. $F'_{0.015}$, therefore there are at most $\binom{d/2}{\gamma}$ ways to choose N'_x.

Combining the bounds from Lemmas 13 and 15 establishes (2) for $j = 2$.