Phase Transitions for Scale-Free SAT Formulas

Tobias Friedrich, Anton Krohmer, Ralf Rothenberger, Andrew M. Sutton
Algorithm Engineering Group
Hasso Plattner Institute
Potsdam, Germany

Abstract

Recently, a number of non-uniform random satisfiability models have been proposed that are closer to practical satisfiability problems in some characteristics. In contrast to uniform random Boolean formulas, scale-free formulas have a variable occurrence distribution that follows a power law. It has been conjectured that such a distribution is a more accurate model for some industrial instances than the uniform random model. Though it seems that there is already an awareness of a threshold phenomenon in such models, there is still a complete picture lacking. In contrast to the uniform model, the critical density threshold does not lie at a single point, but instead exhibits a functional dependency on the power-law exponent. For scale-free formulas with clauses of length $k = 2$, we give a lower bound on the phase transition threshold as a function of the scaling parameter. We also perform computational studies that suggest our bound is tight and investigate the critical density for formulas with higher clause lengths. Similar to the uniform model, on formulas with $k \geqslant 3$, we find that the phase transition regime corresponds to a set of formulas that are difficult to solve by backtracking search.

Introduction

Despite the worst-case complexity results of SAT, many large industrial instances can be solved efficiently by modern solvers. These instances, unlike formulas generated uniformly at random, appear to have some kind of underlying tractable substructure that can be exploited by solvers during preprocessing or execution (Williams, Gomes, and Selman 2003; Kullmann 2004; Ansótegui et al. 2008). Nevertheless, most theoretical work on SAT instance distributions has focused almost exclusively on the *uniform random* distribution. Uniform random formulas are easy to construct, and are comparatively more accessible to probabilistic analysis due to their statistical uniformity.

A narrow focus on uniform random instances comes with a risk of driving SAT research in the wrong direction (Kautz and Selman 2007) because such instances do not possess the same structural properties as those encountered in practice. It is also well-known that solvers that have been tuned to perform well on one class of instances do not necessarily perform well on another (Birattari 2009). To address this, the SAT community has expanded their view to study so-called *industrial* problem instances. Industrial instances arise from problems in practice, such as hardware and software verification, automated planning and scheduling, and circuit design. Empirically, industrial instances appear to have strongly different properties than formulas generated uniformly at random, and SAT solvers behave very differently when applied to them (Crawford and Baker 1994; Konolige 1994). However, benchmark industrial instances are typically only available on an instance-by-instance basis, so they are impossible to generate for any specific given parameter setting such as problem size; and empirical results are difficult to generalize (Rish and Dechter 2000). Moreover, rigorous theoretical results have been out of reach for such problems because a formal definition of their distribution is still lacking (Williams, Gomes, and Selman 2003). In fact, generating synthetic instances that are more similar to real-world instances is one of the "Ten Grand Challenges" proposed in satisfiability research over the last two decades (Selman, Kautz, and McAllester 1997; Selman 2000; Kautz and Selman 2003; 2007).

An observable difference between uniform random instances and real world instances is the statistics of variable occurrence. On the uniform random model, the occurrence of variables is strongly concentrated around its expectation. On the other hand, very large variations in variable occurrence can be observed on real world instances such as those that arise in bounded model checking and software verification. It has been conjectured (Boufkhad et al. 2005) that this property might be modeled well by a random formula model with a power-law variable distribution. In such a distribution, the fraction of variables that occur z times in a formula is proportional to $z^{-\beta}$, where the constant β is called the *power-law exponent*. The left plot of Figure 1 shows the correspondence between the variable distributions of industrial formulas and random power-law formulas. The figure is generated by plotting the empirical cumulative variable distribution of two groups of industrial formulas selected from SAT Race 2015 competition: `hwmcc10` (hardware model checking) and `SAT_dat` (IBM formal verification suite). We used the Clauset-Shalizi-Newman method (2009) to estimate the power-law exponent in the industrial data and match it in the randomly generated formula. A more detailed look at the variable distribution statistics of industrial formulas is

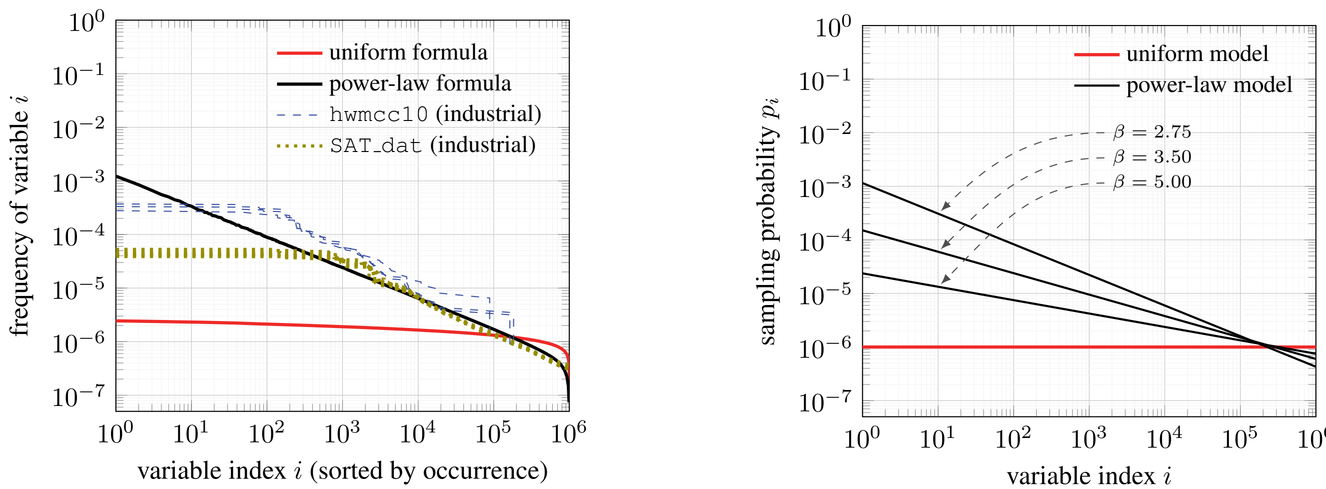

Figure 1: Empirical variable occurrence frequency (**left**) for formulas in two industrial categories from SAT Race 2015. `hwmcc10` contains three formulas with $8 \times 10^4 < n < 2 \times 10^5$; `SAT_dat` contains five formulas with $10^5 < n < 1.5 \times 10^6$. The industrial formulas are compared to a random power-law formula ($n = 10^6$, $m = 4.5 \times 10^6$, $\beta = 2.75$) and a uniform random formula ($n = 10^6$, $m = 4.5 \times 10^6$). The variable frequencies of the industrial instances are closer to the random power-law formula than to the uniform random formula. In fact, a vast majority of the industrial variable frequencies match very closely to the powerlaw distribution. Variable occurrence probability p_i (**right**) as a function of variable index i as described in Equation (2) for $n = 10^6$, $\beta = 2.5, 3.5, 5.0$ scale-free distributions and uniform distribution.

provided by Ansótegui, Bonet, and Levy (2009a).

The *constraint graph* of a formula Φ is a graph $G = (V, E)$ where V corresponds to the variables of Φ and $(u, v) \in E$ iff u and v appear together in a clause. The power-law variable distribution gives rise to an underlying constraint graph with the so-called *scale-free* property. The presence of this scale-free property on real-world formulas could explain the success of complete solvers on industrial instances (Ansótegui, Bonet, and Levy 2009a), since it has been shown that they typically have a small set of key variables that, when set correctly, make the formula easy (Williams, Gomes, and Selman 2003). Others have further conjectured that even incomplete stochastic local search algorithms could also exploit scale-free structure in SAT formulas (Tompkins and Hoos 2010).

To address the need for industrial-like models and motivated by the conjectured scale-free property in real world problems noted by Boufkhad et al. (2005), Ansótegui et al. (2009a; 2009b) proposed a number of non-uniform SAT distributions including a model in which variables are selected from a power-law distribution. The resulting instances have the benefit that they can be generated at random, but they also bear the claimed computational properties observed in practical applications (Ansótegui, Bonet, and Levy 2009b).

The study of random SAT formulas is concerned with characterizing probability distributions over satisfiability formulas. The *constraint density* of a distribution of formulas on n variables and m clauses is measured as the ratio of clauses to variables m/n. A *phase transition* in a random satisfiability model is the phenomenon of a sharp transition as a function of constraint density between formulas that are almost surely satisfiable and formulas that are almost surely not satisfiable. The location of such a transition is called the *critical density*. Ansótegui et al. (2009b) experimentally identified a "phase transition point" at a constant density for the scale-free model. They found that the critical density appears to depend on the power-law exponent of the variable occurrence distribution. However, its precise location and functional dependence on both the power-law exponent and density is not currently known. Furthermore, so far no rigorous bounds exist for formulas of any clause length.

Our contribution. In this paper, we specifically address the phase transition phenomenon in scale-free k-SAT formulas. We sketch a proof for a lower bound on the location of the threshold for the case of $k = 2$. Our bound implies that the critical density scales as a simple polynomial function of the power-law exponent β of the formula. We present empirical results that suggest our bound is tight. For $k > 2$, we perform a computational study to provide a more detailed picture of the phase transition for scale-free formulas, and clarify the regime of hard formulas. We conjecture that the critical density also scales as some polynomial in β for $k > 2$, and we also give evidence for a dependence on k.

Similar to the uniform model, we find that the threshold corresponds to hard formulas. In the scale-free model, this corresponds to a region of hard formulas that depends both on the density and the power-law exponent. We perform a short empirical study that suggests this hard region is relatively consistent across different types of solvers.

Related work on uniform random SAT. In the *uniform random distribution* each k-CNF formula over n variables and m clauses occurs with equal probability. This distribution was popularized in a seminal paper by Mitchell, Selman and Levesque (1992) who were interested in whether hard

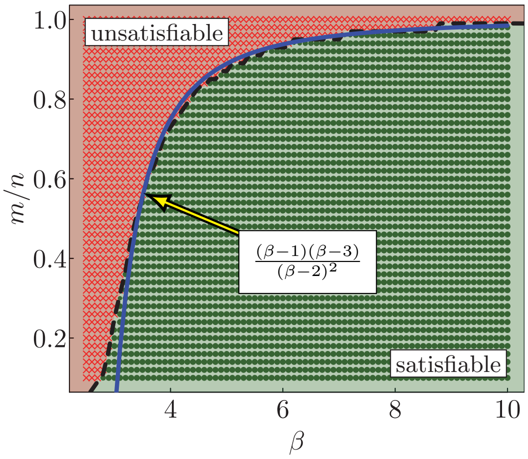

Figure 2: Phase diagram for scale-free 2-SAT formulas with $n = 10^7$ variables. We empirically observe a sharp phase transition (**– –**), which closely matches the theoretical bound of Theorem 1 (**—**).

instances of SAT come up often, or rather if they are always a result of encodings tailored to a specific purpose. They also showed that a backtracking complete solver exhibited an empirical easy-hard-easy pattern at the transition. The explanation for this pattern is that formulas near the critical density have few satisfying assignments (unlike low-density formulas), but still many variables must be assigned before finding a solution or deriving a contradiction (unlike high-density formulas) resulting in deep search trees.

The study of a threshold phenomenon in the uniform model has been the focus of intense study in the last two decades. The *satisfiability threshold conjecture* asserts if Φ is a formula drawn uniformly at random from the set of all k-CNF formulas with n variables and m clauses, there exists a real number r_k such that $\lim_{n\to\infty} \Pr\{\Phi \text{ is satisfiable}\}$ vanishes for $m/n < r_k$, and is equal to one for $m/n < r_k$. Friedgut (1999) showed that the transition is sharp, even though its location is not known exactly for all values of k (and may also depend on n). For $k = 2$, the critical threshold is $r_2 = 1$ (Chvátal and Reed 1992). For $k \geqslant 3$, recently Coja-Oghlan (2014) proved that $r_k = 2^k \ln 2 - \frac{1}{2}(1 + \ln 2) + o(1)$ where the asymptotic term vanishes as $k \to \infty$. Ding et al. (2015) derived an exact representation of the threshold for all $k \geqslant k_0$, where k_0 is a large enough constant. Gent and Walsh (1996; 1999) performed a detailed experimental study on the phase transition for random k-SAT and several other models. For random k-SAT, they found the hardest problems for backtracking solvers are associated with a constraint gap that induces hard unsatisfiable branches in the search tree.

Related work on non-uniform random SAT. The study of phase transitions has been extended to other random distributions. Cooper et al. (2007) identified the existence of a sharp phase-transition in random 2-SAT formulas with prescribed literal occurrences. More specifically, if all literals appear equally often, we obtain regular random k-CNF formulas. This model is claimed to be more difficult to solve than the uniform model because solvers cannot exploit variations in variable occurrence. Boufkhad et al. (2005) used the result of Cooper et al. (2007) to show that the regular random k-CNF distribution exhibits a sharp phase transition at $m/n = 1$ for $k = 2$, and established upper and lower bounds on the critical density for $k = 3$. Their results suggest that the transition of the regular model occurs at much lower densities than with the uniform model.

Giráldez-Cru and Levy (2015) argue that industrial instances exhibit *modularity* and propose a random instance generator that partitions the variables into c disjoint sets of size n/c. Clauses are then chosen uniformly at random so that each clause contains literals only from the same community with probability p, otherwise with probability $1 - p$ they are chosen uniformly out of the set of variables such that each variable comes from a different community. They argue that when the communities do not overlap ($p = 1$), and their size tend to infinity, the phase transition point is the same for uniform random formulas.

Random scale-free formulas

We consider random k-SAT formulas Φ on n variables and m clauses. We denote by $x_1, \ldots, x_n$ the Boolean variables. A clause of size k is a disjunction of exactly k literals $\ell_1 \vee \ldots \vee \ell_k$, where each literal assumes a (possibly negated) variable. Following conventions, we write $|\ell|$ to denote the variable corresponding to literal ℓ. Finally, a formula Φ in conjunctive normal form is a conjunction of clauses $c_1 \wedge \ldots \wedge c_m$. We say that Φ is satisfiable if there exists an assignment of its variables $x_1, \ldots, x_n$ such that the formula evaluates to 1.

To construct a random satisfiability formula in the scale-free model, we sample each clause independently at random. In contrast to the classical uniform random model, however, the probabilities $p_i := \Pr(X = x_i)$ to choose a variable x_i are non-uniform. In particular, a scale-free formula is generated by using a *power-law* distribution for the variable distribution. To this end, we assign each variable x_i a *weight* w_i and sample it with probability

$$p_i := \Pr(X = x_i) = \frac{w_i}{\sum_j w_j}.$$

To achieve a power-law distribution, we assume the concrete weight sequence

$$w_i := \tfrac{\beta-2}{\beta-1}\left(\tfrac{n}{i}\right)^{\frac{1}{\beta-1}} \tag{1}$$

for $i = 1, 2 \ldots, n$, which is a canonical choice for power-law weights, cf. (Chung and Lu 2002a). This sequence guarantees $\sum_j w_j \to n$ for $n \to \infty$ and therefore

$$p_i \to \tfrac{1}{n}\tfrac{\beta-2}{\beta-1}\left(\tfrac{n}{i}\right)^{\frac{1}{\beta-1}}. \tag{2}$$

We plot Equation (2) on the right of Figure 1 for various values of β compared with the flat uniform distribution. To sample Φ, we generate each clause c as follows.

1. Sample k variables independently at random according to the distribution p_i. Repeat until no variables coincide.
2. Negate each of the k variables independently at random with probability $1/2$.

This model for constructing propositional satisfiability formulas first appeared in (Ansótegui, Bonet, and Levy 2009b). Note, however, that the authors of that paper use α instead of β as the power-law exponent and define $\beta = 1/(\alpha-1)$. We instead define β to follow the notational convention of Chung and Lu, cf. (Aiello, Chung, and Lu 2000; Chung and Lu 2002b; 2002a).

Analysis of scale-free 2-SAT

2-SAT is polynomial-time solvable, but still exhibits some typical properties of k-SAT such as a phase transition. 2-SAT was also the first uniform random SAT distribution for which a sharp threshold could be proven (Chvátal and Reed 1992). We show a proof sketch that power-law 2-SAT formulas are satisfiable if the constraint density is small enough. Denote $[n] := \{1, \ldots, n\}$.

Lemma 1. *Whenever $\beta > 3$, it holds for all $1 \leqslant i \leqslant n$*

$$\sum_{\substack{S\subseteq[n]\\|S|=t}} \prod_{i\in S} p_i^2 \leqslant \exp\left(\tfrac{t\cdot(\beta-3)}{n\cdot(\beta-1)}\right) \cdot \left(\tfrac{(\beta-2)^2}{(\beta-1)(\beta-3)}\right)^t \cdot \frac{1}{t!} \cdot n^{-t}.$$

Proof Sketch. The idea is to arrange the elements of the set $S \subseteq [n]$ in decreasing order $s_1 > s_2 > \ldots > s_t$, yielding

$$\sum_{\substack{S\subseteq[n]\\|S|=t}} \prod_{i\in S} p_i^2 = \sum_{s_1=t}^{n} \left(p_{s_1}^2 \sum_{s_2=t-1}^{s_1-1} \left(p_{s_2}^2 \cdots \sum_{s_t=1}^{s_{t-1}-1} p_{s_t}^2\right)\right).$$

This can be estimated inductively, beginning with the innermost sum. In the induction step, the resulting sums are upper bounded by integrals to derive the desired result. □

Theorem 1. *Scale-free random 2-SAT with power-law exponent $\beta > 3$ and clause-variable ratio $m/n < \frac{(\beta-1)(\beta-3)}{(\beta-2)^2}$ is satisfiable with probability $1 - o(1)$.*

Proof Sketch. The proof of this theorem is oriented along the lines of the proof of Theorem 3 from (Chvátal and Reed 1992). We define a *bicycle* of length t to be a sequence of $t+1$ clauses of the form $(u, \ell_1), (\bar{\ell}_1, \ell_2), \ldots, (\bar{\ell}_{t-1}, \ell_t), (\bar{\ell}_t, v)$, where $\ell_1, \ldots, \ell_t$ are literals of distinct variables and $u, v \in \{\ell_1, \ldots, \ell_t, \bar{\ell}_1, \ldots, \bar{\ell}_t\}$. Chvátal and Reed (1992, Theorem 3) show that every unsatisfiable 2-CNF contains a bicycle. To upper bound the probability that Φ is unsatisfiable it is therefore sufficient to upper bound the probability that Φ contains a bicycle. To do so, we calculate the expected number of bicycles and use Markov's inequality.

Observe that the probability to sample a clause $(\ell \vee h)$ is exactly

$$\Pr(\ell \vee h) := \frac{p_{|\ell|} \cdot p_{|h|}}{2\left(1 - \sum_{i=1}^{n} p_i^2\right)} = \frac{C}{2} \cdot p_{|\ell|} \cdot p_{|h|}, \qquad (3)$$

since a pair of different variables is sampled with probability $2 \cdot p_{|\ell|} \cdot p_{|h|}$; there are four combinations of signs; and the probability that no variable appears twice is $C^{-1} = \left(1 - \sum_{i=1}^{n} p_i^2\right) = 1 - \Theta(\frac{1}{n})$ for our given weight sequence whenever $\beta > 3$.

Now let X denote the random variable counting the number of bicycles that appear in F. One can show combinatorially that

$$\mathbb{E}[X] \leqslant \sum_{t=2}^{n} 2 \cdot m^{t+1} \cdot t! \cdot C^{t+1} \cdot t^2 p_{\max}^2 \sum_{\substack{S\subseteq[n]\\|S|=t}} \prod_{i\in S} p_i^2.$$

By Lemma 1 and using $r := m/n$, $p_{\max} = \Theta(n^{-\frac{\beta-2}{\beta-1}})$, and $C = 1 + \Theta\left(\frac{1}{n}\right)$, we obtain that the right-hand side is at most

$$\begin{aligned} &2 \cdot e^{\frac{\beta-3}{\beta-1}} \cdot r \cdot n \cdot C \sum_{t=2}^{n} \cdot t^2 p_{\max}^2 \left(Cr \tfrac{(\beta-2)^2}{(\beta-1)(\beta-3)}\right)^t \\ &= \mathcal{O}\left(n^{\frac{3-\beta}{\beta-1}}\right) \sum_{t=2}^{n} t^2 \left(r \tfrac{(\beta-2)^2}{(\beta-1)(\beta-3)}\right)^t. \end{aligned} \qquad (4)$$

Since we assume $r\frac{(\beta-2)^2}{(\beta-1)(\beta-3)} < 1$, there is a $t_0 = \operatorname{polylog}(n)$ with $t_0^2 \cdot \left(r\frac{(\beta-2)^2}{(\beta-1)(\beta-3)}\right)^{t_0} < n^{-2}$, i.e.

$$\sum_{t=2}^{n} t^2 \left(r \tfrac{(\beta-2)^2}{(\beta-1)(\beta-3)}\right)^t \leqslant t_0^2 + n^{-1} = \mathcal{O}(\operatorname{polylog}(n)).$$

This gives us $\mathbb{E}[X] = \mathcal{O}(n^{\frac{3-\beta}{\beta-1}} \cdot \operatorname{polylog}(n)) = o(1)$, since we require $\beta > 3$. □

This yields a lower bound for the critical density of scale-free random 2-SAT formulas. Figure 2 shows that this theoretical bound matches our empirical observations.

Empirical study of scale-free k-SAT

In this section, we experimentally investigate the behavior of scale-free k-SAT. We report the satisfiability of random formulas depending on the constraint density m/n and the power-law exponent β in Figure 2 (for $k = 2$) and Figure 3 (for $k = 3, 4$). The phase diagram in Figure 2 and in the leftmost panels of Figure 3 are generated as follows. In an interest to eliminate statistical fluctuations that can arise at small problem sizes, we set n very large, specifically $n = 10^7$ for $k = 2$ and $n = 10^6$ for $k > 2$. For the appropriate value of n, 50 scale-free formulas are generated for each $\beta = 1.5, 1.6, \ldots, 3.5$ and each m such that $m/n = 1/10, 2/10, \ldots, 5$ (for $k = 2$ we use increments of $1/50$ with maximum $m/n = 1$). Each resulting formula is solved by the CDCL-based solver `MapleCOMSPS` (Liang et al. 2016) with a time cutoff of 900 seconds (15 min).

If the satisfiability of the formula cannot be determined within this time, the formula is marked as "hard", and its satisfiability state is unknown. On the phase diagrams, each point corresponds to a set of 50 formulas at a given density and scale parameter. If all 50 formulas were unsatisfiable, a red cross (×) is drawn. If some formulas are satisfiable, a green dot (●) is drawn with the size of the dot corresponding to the fraction of satisfiable instances. Note that the threshold appears to be sharp, and in most cases, either all formulas

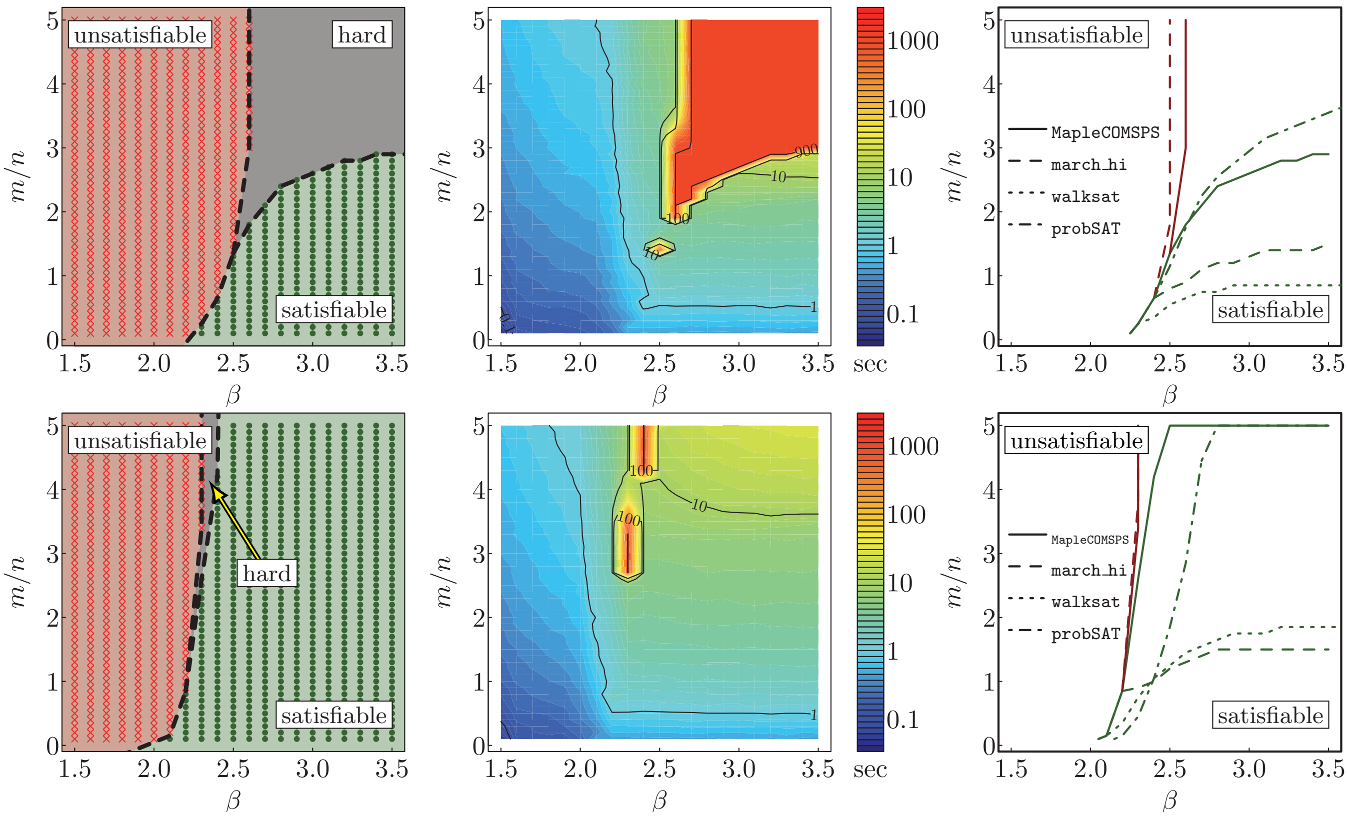

Figure 3: Phase diagrams (**left**) and timing contour plots (**center**) for scale-free k-SAT with $n = 10^6$ and $k = 3$ (**top**) and $k = 4$ (**bottom**). The phase diagrams show a clear phase transition from unsatisfiable (✕) to satisfiable (●) and a patch of very hard instances (■), close to the phase transition for higher m/n and β. The contour plots and heat maps in the center column report mean solver time on the formulas (blue=fast,red=slow); solver run time strongly increases around the phase transition. Comparison of threshold bounds (**right**) proposed by four different solvers. As a function of β: the upper bound on density for the unsatisfiable phase is drawn in red; lower bound on density for satisfiable phase drawn in green.

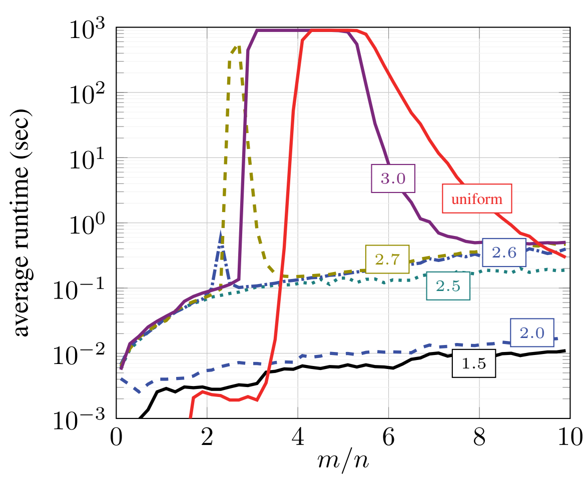

Figure 4: Solver time (MiniSAT) for different values of β depending on density m/n. Solver cutoff: 900s; scale-free parameters: $n = 10^4$, $k = 3$; uniform parameters: $n = 500$, $k = 3$. Uniform formulas and scale-free formulas with $\beta > 2.6$ show the typical easy-hard-easy pattern. Small uniform instances are still harder than all shown scale-free instances for $m/n \in [5, 9]$.

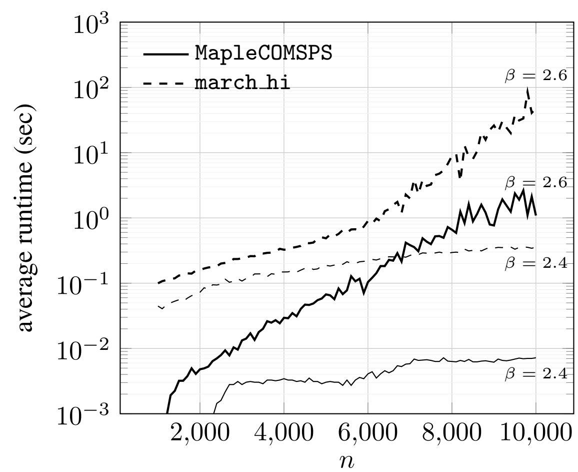

Figure 5: Scaling behavior (MapleCOMSPS, march_hi) at fixed density $m/n = 2.28$ and $k = 3$ at critical point ($\beta = 2.6$) and below critical point ($\beta = 2.4$). Both solvers scale exponentially at the critical point (corresponding to the lower left corner of "hard" region in top left phase diagram in Figure 3). Slightly below the critical point, both solvers scale more efficiently.

were satisfiable, all were unsatisfiable, or all were hard. Moreover, the sat/unsat transition appears to shift to lower overall β values when moving from $k = 3$ to $k = 4$.

To understand the effect of the model parameters on solver time, we measured the average time to determine the satisfiability of all formulas at each $(\beta, m/n)$ point. When $k = 2$, solver time was not generally affected by the threshold, though there is an expected (polynomial) dependency on formula size. We omit the details for this case due to space constraints. For $k = 3, 4$ we report the mean solver time (in seconds) in a contour plot in the center panel of Figure 3.

Our solver choice is motivated by the good performance of `MapleCOMSPS` on the Application Benchmarks at the SAT 2016 competition. `MapleCOMSPS` implements machine-learning branching heuristics and is based on `MiniSAT` (Eén and Sörensson 2003), which has also performed well in industrial categories in previous SAT competitions (Le Berre and Simon 2006). For $k = 3$, there is an expanding region of hard formulas that begins at densities above 2 and power-law exponent $\beta \geqslant 2.5$. For $k = 4$, we see a crossover at around 2.33, and the hard region seems to lie at higher densities.

In order to determine how other types of solvers behave in the model, we also repeated the experiment with `march_hi` (Heule and van Maaren 2009), a DPLL-based solver employing look-ahead heuristics to select branching variables, and two stochastic local search (SLS) solvers, `WalkSAT` (Selman, Kautz, and Cohen 1994) and `probSAT` (Balint and Schöning 2014). We plot the upper and lower bounds of the sat/unsat threshold proposed by all solvers for $k = 3, 4$ in the right-most panel of Figure 3. The lines can be interpreted as follows. At each β value, the highest (resp., lowest) density at which the majority of formulas are successfully determined to be unsatisfiable (resp., satisfiable) yields a proposed upper bound (resp., lower bound) on the threshold at that β. The upper bounds (at the unsat region) are drawn in red and the lower bounds (at the sat region) are drawn in green. Note that the SLS solvers are incomplete in the sense that they can only decide satisfiability. Therefore, they can only propose lower bounds on the threshold.

Interestingly, the two backtracking solvers behave similarly in the unsatisfiable phase, but `march_hi` fails for much lower densities in the satisfiable phase. Its performance in the satisfiable phase is similar to that of `WalkSAT`, though even `WalkSAT` solves slightly higher densities at large β in the $k = 4$ model. Furthermore, despite the overall weak performance of `WalkSAT`, the other SLS solver (`ProbSAT`) even achieves the best proposed lower bound for $k = 3$.

Figure 4 investigates easy-hard-easy patterns at fixed β values. The easy-hard-easy pattern is a well-known phenomenon for backtracking solvers in the uniform model. Difficult formulas appear at the critical density and become easier again as density is increased. This can be an overly simplistic perspective. Coarfa et al. (2003) summarize this by pointing out that the pattern is observed only when fixing n and varying m/n. In this case the class of instances is finite. On the other hand, fixing density and varying n reveals a different "slice" of the picture. For example, Chvátal and Szemerédi (1988) showed that any resolution proof of a high density uniform random formula (e.g., $m/n \geqslant 5.8$ for $k = 3$) must generate a number of clauses that grows exponentially with n.

Nevertheless, understanding hardness as a function of density at fixed orders can help to give a picture of where the troubling regions are for solvers at finite n, and we observe in Figure 4 that such a pattern emerges as β increases. The densities at which the hard formulas emerge also vary with β. The solver cutoff is 900s, and so the peaks appear truncated. Such a relatively low cutoff means high-density uniform random formulas only become easy again for $n = 500$ (see above point). On the other hand, for scale-free formulas, we are able to observe the pattern even for $n = 10^4$. This supports the conjecture that scale-free random formulas are easier than uniform random formulas for backtracking solvers.

Finally in Figure 5 we observe in a semi-log plot the scaling of mean solver time as a function of n at the critical point $\beta = 2.6$, $m/n = 2.28$ in the $k = 3$ model (lower left of the hard region in the top left of Figure 3). At this point (for $n \leqslant 10^4$), we find that roughly half of the formulas are unsatisfiable. The solver times scale exponentially, but with seemingly similar bases. Reducing the exponent only slightly ($\beta = 2.4$) results in significantly more efficient scaling.

Conclusion

The characterization of the satisfiability threshold on uniform random formulas has been an important and challenging research program spanning several decades. In this paper, we open a new line of inquiry into the satisfiability threshold of *scale-free* k-SAT: a distribution that better models the variable occurrence statistics of many real-world problems. We give a lower bound on the location of the threshold for $k = 2$, and present an empirical picture of the threshold that is described by a line in the plane determined by formula density and power-law exponent.

Hard instances of many random combinatorial problem instances can often be found near the critical values of order parameters (Cheeseman, Kanefsky, and Taylor 1991). In scale-free k-SAT, we also find the hardest instances appear in an expanding region that emerges along the threshold line. Formulas with low power-law exponent β are easy to refute, most likely because they contain simple contradictions. On the other hand, those with small constraint density but sufficiently large β are typically underconstrained and contradiction-free. These formulas are easily solved by both complete backtracking and SLS solvers. On formulas that cluster along the phase transition, we conjecture that, similar to the uniform model (Gent and Walsh 1996) backtracking solvers make spurious decisions early in the search tree. Moreover, solutions become more sparsely distributed in the search space, rendering SLS solvers ineffective.

In general, the study of non-uniform random problem distributions can provide better insights into the instance structure of practical applications. Trajectories for future work include developing a matching upper bound for 2-SAT and generalizing the threshold bounds to $k > 2$. It is also interesting to consider non-uniform random instances of other problems that undergo a phase transition such as graph coloring (Cheeseman, Kanefsky, and Taylor 1991) and vertex cover (Hartmann and Weigt 2001).

References

Aiello, W.; Chung, F.; and Lu, L. 2000. A random graph model for massive graphs. In *Proceedings of 32nd STOC*, 171–180.

Ansótegui, C.; Bonet, M. L.; Levy, J.; and Manyà, F. 2008. Measuring the hardness of SAT instances. In *Proceedings of 23rd AAAI*, 222–228.

Ansótegui, C.; Bonet, M. L.; and Levy, J. 2009a. On the structure of industrial SAT instances. In *Proceedings of 15th CP*, 127–141.

Ansótegui, C.; Bonet, M. L.; and Levy, J. 2009b. Towards industrial-like random SAT instances. In *Proceedings of 21st IJCAI*, 387–392.

Balint, A., and Schöning, U. 2014. probSAT and pprobSAT. In *SAT Competition 2014*.

Birattari, M. 2009. *Tuning Metaheuristics: A Machine Learning Perspective*, volume 197 of *Studies in Computational Intelligence*. Springer.

Boufkhad, Y.; Dubois, O.; Interian, Y.; and Selman, B. 2005. Regular random k-SAT: Properties of balanced formulas. *Journal of Automated Reasoning* 35(1-3):181–200.

Cheeseman, P.; Kanefsky, B.; and Taylor, W. M. 1991. Where the really hard problems are. In *Proceedings of 12th IJCAI*, 331–340.

Chung, F., and Lu, L. 2002a. The average distances in random graphs with given expected degrees. *Proceedings of the National Academy of Sciences* 99(25):15879–15882.

Chung, F., and Lu, L. 2002b. Connected components in random graphs with given expected degree sequences. *Annals of Combinatorics* 6(2):125–145.

Chvátal, V., and Reed, B. 1992. Mick gets some (the odds are on his side). In *Proceedings of 33rd FOCS*, 620–627.

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

Clauset, A.; Shalizi, C. R.; and Newman, M. E. J. 2009. Power-law distributions in empirical data. *SIAM Review* 51(4):661–703.

Coarfa, C.; Demopoulos, D. D.; Aguirre, A. S. M.; Subramanian, D.; and Vardi, M. Y. 2003. Random 3-SAT: The plot thickens. *Constraints* 8(3):243–261.

Coja-Oghlan, A. 2014. The asymptotic k-SAT threshold. In *Proceedings of 46th STOC*, 804–813.

Cooper, C.; Frieze, A.; and Sorkin, G. B. 2007. Random 2-SAT with prescribed literal degrees. *Algorithmica* 48(3):249–265.

Crawford, J. M., and Baker, A. B. 1994. Experimental results on the application of satisfiability algorithms to scheduling problems. In *Proceedings of 12th AAAI*, 1092–1097.

Ding, J.; Sly, A.; and Sun, N. 2015. Proof of the satisfiability conjecture for large k. In *Proceedings of 47th STOC*, 59–68.

Eén, N., and Sörensson, N. 2003. An extensible SAT-solver. In *Proceedings of 6th SAT*, 502–518.

Friedgut, E. 1999. Sharp thresholds of graph properties, and the k-SAT problem. *Journal of the American Mathematical Society* 12(4):1017–1054.

Gent, I. P., and Walsh, T. 1996. The satisfiability constraint gap. *Artificial Intelligence* 81(1-2):59–80.

Gent, I. P., and Walsh, T. 1999. Beyond NP: the QSAT phase transition. In *Proceedings of 16th AAAI*, 648–653.

Giráldez-Cru, J., and Levy, J. 2015. A modularity-based random SAT instances generator. In *Proceedings of 24th IJCAI*, 1952–1958.

Hartmann, A. K., and Weigt, M. 2001. Statistical mechanics perspective on the phase transition in vertex covering of finite-connectivity random graphs. *Theoretical Computer Science* 265(1-2):199–225.

Heule, M., and van Maaren, H. 2009. march_hi. In *SAT Competition 2009*.

Kautz, H., and Selman, B. 2003. Ten challenges redux: Recent progress in propositional reasoning and search. In *Proceedings of 9th CP*, 1–18.

Kautz, H., and Selman, B. 2007. The state of SAT. *Discrete Applied Mathematics* 155(12):1514–1524.

Konolige, K. 1994. Easy to be hard: Difficult problems for greedy algorithms. In *Proceedings of 4th KR*, 374–378.

Kullmann, O. 2004. Upper and lower bounds on the complexity of generalised resolution and generalised constraint satisfaction problems. *Annals of Mathematics and Artificial Intelligence* 40(3-4):303–352.

Le Berre, D., and Simon, L. 2006. Preface to the special volume on the SAT 2005 competitions and evaluations. *Journal on Satisfiability, Boolean Modeling and Computation*.

Liang, J. H.; Oh, C.; Ganesh, V.; Czarnecki, K.; and Poupar, P. 2016. MapleCOMSPS, MapleCOMSPS_lrb, MapleCOMSPS_CHB. In *SAT Competition 2016: Solver and Benchmark Descriptions*, volume B-2016-1 of *Department of Computer Science Series of Publications B, University of Helsinki*, 52–53.

Mitchell, D. G.; Selman, B.; and Levesque, H. J. 1992. Hard and easy distributions of SAT problems. In *Proceedings of 10th AAAI*, 459–465.

Rish, I., and Dechter, R. 2000. Resolution versus search: Two strategies for SAT. *Journal of Automated Reasoning* 24(1-2):225–275.

Selman, B.; Kautz, H. A.; and Cohen, B. 1994. Noise strategies for improving local search. In *Proceedings of 12th AAAI*, 337–343.

Selman, B.; Kautz, H.; and McAllester, D. 1997. Ten challenges in propositional reasoning and search. In *Proceedings of 15th IJCAI*, 50–54.

Selman, B. 2000. Satisfiability testing: recent developments and challenge problems. In *Proceedings of 15th LICS*, 178–178.

Tompkins, D. A., and Hoos, H. H. 2010. Dynamic scoring functions with variable expressions: New SLS methods for solving SAT. In *Proceedings of 13th SAT*, 278–292.

Williams, R.; Gomes, C.; and Selman, B. 2003. Backdoors to typical case complexity. In *Proceedings of 18th IJCAI*, 1173–1178.