Computational Complexity and Phase Transitions (extended abstract*)

Gabriel Istrate
Center for Nonlinear Studies, CIC-3 Division and TSA-SA
Los Alamos National Laboratory, Mail Stop B258,
Los Alamos, NM 87545, U.S.A.
e-mail: istrate@lanl.gov

Abstract

Phase transitions in combinatorial problems have recently been shown [2] to be useful in locating "hard" instances of combinatorial problems. The connection between computational complexity and the existence of phase transitions has been addressed in Statistical Mechanics [2] and Artificial Intelligence [3], but not studied rigorously.

We take a first step in this direction by investigating the existence of sharp thresholds for the class of *generalized satisfiability problems*, defined by Schaefer [4]. In the case when all constraints have a special clausal form we completely characterize the generalized satisfiability problems that have a sharp threshold. While NP-completeness does *not* imply the sharpness of the threshold, our result suggests that the class of counterexamples is rather limited, as all such counterexamples can be predicted, with constant success probability by a *single* procedure.

1 Introduction

Which combinatorial problems have "hard" instances? Computational Complexity is the main theory that attempts to provide answers to this question. But it is not the only one. While the concept of NP-complete problem, as a paradigm for "problem with hard instances", has permeated a wide range of fields, from Computational Biology to Economics, it is not usually considered extremely relevant by practitioners. This happens because NP-completeness is an overly pessimistic, worst-case, concept, and in fact if we're not really careful about the random model, "most" instances of many NP-complete problems turn out to be "easy".

Much insight in locating the regions "where the really hard instances are" has come from an analogy with Statistical Mechanics, in the context of *phase transitions in combinatorial problems.* Recent studies [2] have shown that a certain type of phase transitions (called *first-order phase transitions*) is responsible for the exponential slowdown of many natural algorithms when run on instances at the transition point.

A natural, and early stated question is whether there exists any connection between computational complexity and the existence of a phase transition. Obtaining an answer to this question is further complicated by the fact that the physicists' and computer scientists' concepts of phase transitions are different: the former pertains to combinatorial optimization, and is called *order-disorder phase transition*, while the latter applies to decision problems and is called *threshold property*, more specifically a restricted form of threshold property called *sharp threshold*[1]. It is this type of phase transitions we're primarily interested in this paper.

*an extended version will be available shortly as [1].

[1]see definition 3.

The above question has been asked for both types of phase transitions: Fu [5] argued that there should be no connection between worst-case computational complexity and the existence of an order–disorder phase transition, by showing that an NP-complete problem, number partition, has no order-disorder phase transition (however see [6] that argues that number partition *has* an order-disorder phase transition under a different random model). The case of decision problems is even more spectacular: in a paper that proved very influential in the Artificial Intelligence community [3], Cheeseman, Kanefsky and Taylor conjectured that roughly the difference between tractable and intractable problems, specifically between problems in P and NP-complete problems is that:

1. NP-complete problems have a phase transition (sharp threshold) with respect to "some" order parameter.

2. in contrast, problems in P lack such a threshold.

Their conjecture was at best wishful thinking. First, they did not make it precise enough, by specifying what an order parameter is. Second, they had no evidence supporting such a radical statement. In fact, examples of problems in P that *do* have a sharp threshold with respect to a "reasonable" order parameter had already long been known (for instance the probability that a random graph has a connected component of at least, say, $n^{3/4}$ vertices, by the classical results of Erdős and Rényi [7]).

A natural question is *whether there is any connection at all between computational complexity and the existence of a sharp threshold at least for problems that possess some "canonical" order parameter.* One restriction that entails the existence of a canonical order parameter is the very one which was used in defining threshold properties: *monotonicity* [8]. Clearly the above-mentioned example shatters the hope of obtaining a version of (2) even for monotonic problems. A quick argument shows that even (1) should fail: in any polynomial degree there exist both monotone problems that have (or do not have) sharp thresholds. The intuitive reason is that the existence of a sharp threshold is a statistical property, that is not affected by modifying a given problem on a set of instances that has zero measure. On the other hand worst-case complexity is sensitive to such changes. The result is formally stated as Proposition 5.1 in the Appendix.

Given the above argument it would seem that the question has been answered, and that no whatsoever connection exists between the two concepts. However the examples constructed in Proposition 5.1 are rather artificial, and the overall proof is reminiscent of Ladner's [9] result on the structure of polynomial degrees: we can construct a set of the desired complexity by starting with a certain base set and "tuning-up" its worst-case complexity on a set that is "small enough" so that this does not affect the other desirable property of the base set, having a sharp/coarse threshold. The question still remains whether the result remains true if we only consider problems with a certain "natural" structure. After all, this is true in the case of computational complexity: Schaefer [4], showed that, when restricted to the class of *generalized satisfiability problems*, the rich structure of polynomial m-degrees derived from Ladner's results simplifies to only two degrees, P and the degree of NP-complete problems, and obtained a full characterization of such problem.

Definition 1 *Let $S = \{R_1, \ldots, R_p\}$, $R_i \subset \{0,1\}^{r_i}$, be a finite set of relations. An* S-formula *in* n variables *is a finite conjunction of* clauses*, i.e. expressions of the type $R_j(x_{j,1}, \ldots, x_{j,r_j})$, with the variables x_j chosen from a fixed set of n variables $x_1, \ldots, x_n$. $SAT(S)$ is the problem of deciding whether an arbitrary S-formula has a satisfying assignment $x_1 \ldots x_n$ (one that makes each clause true).*

A pleasant feature of Schaefer's framework is that every problem $SAT(S)$ is monotonic. Clearly, an analog of (2) fails in this case as well: the density result Proposition 5.1 is still true for one of the two polynomial degrees, P, as 2-SAT has a sharp threshold [10], while e.g. at-most-2-HORN-SAT has a coarse threshold [11]. On the other hand there exists some evidence that some notion of computational intractability implies the existence of a sharp threshold: in his celebrated result on sharp thresholds for

3-SAT Friedgut gives an example of a NP-complete graph problem having a coarse threshold: the property of containing either a triangle or a "large" clique. From a probabilistic standpoint the second part is "not important". Moreover, his characterization theorem implies that *any* graph theoretic property that fails to have a sharp threshold can be well "approximated" by a tractable property, the property of containing a copy of a fixed graph. Finally, there is an altogether different reasons for a rigorous study of sharp thresholds in satisfiability problems: in this case the notion of a first-order phase transition (that, as mentioned *does have* significant algorithmic implications) has a nice combinatorial interpretation, as a "sudden jump" in the relative size of a combinatorial parameter called *backbone* (see e.g. [14] for definition and discussion). It is easy to show (this is an argument implicitly made in [2], that will be presented in the full version of the paper) that the discontinuity of the backbone implies the existence of a sharp threshold. Therefore studying problems with sharp thresholds is a useful first step towards identifying all satisfiability problems having a first order phase transition.

It is, perhaps, tempting to conjecture that, when restricted to Schaefer's framework an analogue of (1) holds:

Hypothesis 1 *Every generalized satisfiability problem $SAT(S)$ that Schaefer's dichotomy theorem [4] identifies as* NP-complete has a sharp threshold.

We further restrict our framework to the case when all constraints in S have a special, clausal form. In this case we obtain a complete characterization of all sets of constraints S for which $SAT(S)$ has a sharp threshold. In a preliminary version of this paper we claimed that for clausal constraints NP-completeness implies the existence of a sharp threshold. Unfortunately this is not true, as the revised version of our result shows. On the other hand, as displayed by Corollary 1, the class of counterexamples is rather limited: they are those NP-complete problems for which satisfiability of a random instance Φ can be predicted with significant success by a very trivial heuristic: if neither 0^n or 1^n are satisfying assignments then return "unsatisfiable". So the lack of a sharp threshold **does** have algorithmic implications, albeit in a probabilistic sense.

2 Preliminaries

We will work in the context of NP–decision problems, a standard concept in Complexity Theory. For a precise definition see, e.g., [12].

Definition 2 *The* NP-decision problem P is *monotonically decreasing* if for every instance x of P and every witness y for x, y is a witness for every instance z obtained by turning some bits of x from 1 to 0. Monotonically increasing problems are defined similarly.

The three main random model from random graph theory, the so-called *constant probability model*, the *counting* and *multiset* model extend directly to NP-decision problems, and are interchangeable under quite liberal conditions. For technical convenience we will use the constant probability model when proving sharp thresholds and the multiset models when dealing with coarse thresholds. The following is a brief review. The multiset model, denoted $\Omega(n,m)$, and which has two integer parameters n, m. A random sample from $\Omega(n,m)$ is obtained by starting with the string $z = 0^n$, choosing (uniformly at random and with repetition) m bits of z, and flipping these bits to one. When n is known, we use $\mu_m(A)$ to refer to the measure of a set A under this random model. The constant probability model denoted $\Omega_p(n)$ has two parameters, an integer n and a real number $p \in [0,1]$. A random sample from $\Omega_p(n)$ is obtained by starting with the string $z = 0^n$ and then flipping the bits of z to one independently with probability p.

Definition 3 *Let P be any monotonically decreasing decision problem under the constant probability model $\Omega_p(n)$. A function $\overline{\theta}$ is a* threshold function *for P if for every function m, defined on the set of admissible instances and taking real values, we have*

1. *if $p(n) = o(\overline{\theta}(n))$ then $\lim_{n\to\infty} \Pr_{x\in\Omega_p(n)}[x \in P] = 1$, and*

2. *if* $p(n) = \omega(\overline{\theta}(n))$ *then* $\lim_{n\to\infty} \Pr_{x\in\Omega_{p}(n)}[x \in P] = 0$.

P has a sharp threshold *if in addition the following property holds:*

3. *For every* $\epsilon > 0$ *define the functions* $p_\epsilon(n), p_{1/2}(n), p_{1-\epsilon}(n)$ *by*

$$\Pr_{x\in\Omega_{p_\epsilon}(n)}[x \in P] = \epsilon\},$$

$$\Pr_{x\in\Omega_{p_{1/2}}(n)}[x \in P] = 1/2\},$$

$$\Pr_{x\in\Omega_{p_{1-\epsilon}}(n)}[x \in P] = 1-\epsilon\}.$$

Then we have

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

If, on the other hand, for some $\epsilon > 0$ *the amount* $\frac{p_{1-\epsilon}(n)-p_\epsilon(n)}{p_{1/2}(n)}$ *is bounded away from 0 as* $n \to \infty$, *P has a* coarse threshold. *These two cases are not exhaustive as the above quantity could in principle oscillate with n. Nevertheless they are so for most "natural" problems.*

Let $f : \mathbf{N} \to \mathbf{R}$. Define $QEMPTY(f)$ to be the probability that the following queuing chain:

$$\begin{cases} Q_0 = 1, \\ Q_{i+1} = Q_i - 1 + \Xi_{i+1}. \end{cases}$$

(where the Ξ_t's are independent Poisson variables with parameter $f(t)$) ever remains without customers.

Definition 4 *Let* $(a, b) \in \mathbf{N} \times \mathbf{N} \setminus (0, 0)$. *Define* $C_{a,b} = \overline{x}_1 \vee \ldots \vee \overline{x}_a \vee x_{a+1} \vee \ldots \vee x_{a+b}$. *Such a relation is called* clausal constraint.

For a set S as in definition 1 let k be the maximum arity of a relation in S. To avoid trivial cases, we assume that $k \geq 2$. For $i = \overline{1, k}$ let p_i be 1 if clause $\overline{x}_1 \vee \ldots \vee \overline{x}_{i-1} \vee x_i \in S$ and 0 otherwise, and let n_i be 1 if clause $\overline{x}_1 \vee \ldots \vee \overline{x}_i \in S$ and 0 otherwise. Define polynomials $P_i(c) = \sum_{j\geq i} \binom{c}{j-i} \cdot p_j$ and $Q_i(c) = \sum_{j\geq i} \binom{c}{j-i} \cdot n_j$. Let $\delta_k = kp_k + n_k$, $N_S = \binom{n}{k} \cdot \delta_k$, and $\alpha = m/N_S$. Finally, let

$$a_0 = \max\{0\} \cup \{a : C_{a,0} \in S\},$$

$$a_{\geq 1} = \max\{0\} \cup \{a : C_{a,b} \in S, b \geq 1\}.$$

b_0 and $b_{\geq 1}$ are defined similarly with respect to the second component.

3 Main result

Recall that a relation is called *0-valid* (1-valid) if it is satisfied by the assignment "all zeros" ("all ones") and *Horn* (negated Horn) if it is equivalent to a Horn (negated Horn) CNF-formula. When S is Horn the number of clauses in S over n variables is $N_S(1 + o(1))$. For a property T we will use "S is T" as a substitute for "every relation in S is T".

Our main result is

Theorem 3.1 *Let S be a finite set of* clausal constraints.

a. *If S is 0-valid or S is 1-valid then the decision problem SAT(S) is trivial.*

b. *If S is (Horn ∪ 0-valid) or S is (negated Horn ∪ 1-valid) then SAT(S) has a coarse threshold.*

c. *Suppose cases a. and b. do not apply. If*

$$(a_{\geq 1} < a_0 \leq b_0) \vee (b_{\geq 1} < b_0 \leq a_0) \vee$$

$$(a_0 = b_0 = \min\{a_{\geq 1}, b_{\geq 1}\})$$

then SAT(S) has a sharp threshold, otherwise SAT(S) has a coarse threshold.

For reasons of space we can do little but present a rather sketchy outline of the proof of Theorem 3.1. A full version will be given in [1]. The following corollary (of the preceding result and its proof) summarizes the intuition that all NP-complete problems with coarse thresholds are "rather trivial".

Corollary 1 *Suppose S is a finite set of clausal constraints. Then SAT(S) has a coarse threshold* exactly *when at least one of the following (non-exclusive) conditions applies.*

```
Program PUR(Φ):
if Φ (contains no positive unit clause)
   return TRUE
else
   choose such a positive unit clause x
   if (Φ contains x̄ as a clause)
      return FALSE
   else
      let Φ′ be the formula
      obtained by setting x to 1
      return PUR(Φ′)
```

Figure 1: Algorithm PUR

1. *S is Horn.*

2. *S is negated Horn.*

3. *$SAT(S)$ is* NP-complete and has the same threshold function as the property "0^n satisfies Φ".

4. $SAT(S)$ is NP-complete and has the same threshold function as the property "1^n satisfies Φ".

Indeed, in the cases 3 and 4 there *a single trivial algorithm*, that declares the formula unsatisfiable if it is not satisfied by any of the two assignments 0^n and 1^n, and which is correct with a constant probability ϵ over the *whole* range of the parameter p (in the constant probability model).

Observation 1 *In the general case there are other (non-clausal) examples of satisfiability problems with a coarse threshold. Let $R(x,y)$ be the relation "$x \neq y$". Then $SAT(\{R\})$ is essentially the 2-coloring problem, which has a coarse threshold.*

4 Proof sketch

b. This part of the proof is constructive. When Φ is Horn we explicitly determine the probability that a random formula Φ is satisfiable, and then use it to argue that the corresponding (Horn $\cup$ 0-valid) cases also have a coarse threshold. The analysis of the Horn cases is similar to the one when S consists of all Horn clauses of length at most k, that was settled in [11], and is accomplished by analyzing PUR, a natural implementation of positive unit resolution, which is complete for Horn satisfiability.

We regard PUR as working in stages, indexed by the number of variables still left unassigned; thus, the stage number decreases as PUR moves on. We say that *formula Φ survives Stage t* if PUR on input Φ does not halt at Stage t or earlier. Let Φ_i be the formula at the beginning of stage i, and let N_i denote the number of its clauses. We will also denote by $P_{i,t}(N_{i,t})$, the number of clauses of Φ_t of size i and containing one (no) positive literal. Define $\Phi^P_{i,t}$ ($\Phi^N_{i,t}$) to be the subformula of Φ_t containing the clauses counted by $P_{i,t}(N_{i,t})$. The analysis proceeds by showing that we can characterize the evolution of PUR on a random formula by a Markov chain, and is based on the following "Uniformity Lemma" from [11], valid in our context as well:

Lemma 4.1 *Suppose that Φ survives up to stage t. Then, conditional on the values $(P_{1,t}, N_{1,t}, \ldots, P_{k,t}, N_{k,t})$, the clauses in $\Phi^P_{1,t}, \Phi^N_{1,t}, \ldots, \Phi^P_{k,t}, \Phi^N_{k,t}$ are chosen uniformly at random and are independent. Also, conditional on the fact that Φ survives stage t as well, the following recurrences hold:*

$$\begin{cases} P_{1,t-1} = P_{1,t} - 1 - \Delta^P_{01,t} + \Delta^P_{12,t}, \\ N_{1,t-1} = N_{1,t} + \Delta^N_{12,t}, \end{cases}$$

and, for $i = \overline{2,k}$,

$$\begin{cases} P_{i,t-1} = P_{i,t} - \Delta^P_{0i,t} - \Delta^P_{(i-1)i,t} + \Delta^P_{i(i+1),t}, \\ N_{i,t-1} = N_{i,t} - \Delta^N_{(i-1)i,t} + \Delta^N_{i(i+1),t}, \end{cases}$$

where

$$\begin{cases} \Delta^P_{01,t} = B(P_{1,t} - 1, 1/t), \\ \Delta^P_{(i-1)i,t} = B(P_{i,t}, (i-1)/t), \\ \Delta^P_{0i,t} = B(P_{i,t} - \Delta^P_{(i-1)i,t}, 1/t), \\ \Delta^N_{(i-1)i,t} = B(N_{i,t}, i/t), \\ \Delta^P_{k(k+1),t} = \Delta^N_{k(k+1),t} = 0. \end{cases}$$

The main intuition for the proof is that with high probability the binomial expressions in the previous formulas are close to their expected values. The proof of this very intuitive statement is conceptually simple, but technically somewhat involved, and mirrors the proof in [11]. So all it remains is to characterize the mean values of $P_{i,t}$, $N_{i,t}$. We only outline the main steps of this computation in the sequel, assuming that the above mentioned concentration results hold. Define $x_{i,t}, y_{i,t}$ by

$$\begin{cases} E[P_{i,t}] = i \cdot \binom{t}{i} \cdot x_{i,t}, \\ E[N_{i,t}] = \binom{t}{i} \cdot y_{i,t}. \end{cases}$$

Then it is easy to see that sequences $x_{i,t}, y_{i,t}$, $i \geq 2$ verify the recurrences:

$$\begin{cases} x_{i,t-1} = x_{i,t} + x_{i+1,t}, \\ y_{i,t-1} = y_{i,t} + y_{i+1,t}. \end{cases}$$

Define the vector sequence $(Z_t)_{t\geq 0} \in \mathbf{R}^{k-1}$ by $Z_{t+1} = A \cdot Z_t$, with $A = (a_{i,j})$,

$$a_{i,j} = \begin{cases} 1, & \text{if } j = i+1, \\ 0, & \text{otherwise.} \end{cases}$$

It is easy to see that both sequences $(x_{i,t})_t$ and $(y_{i,t})_t$ satisfy the same recurrence as Z_t. A simple computation shows that $A^k_{i,j} = \binom{k}{j-i}$ (where, for $t < 0$, $\binom{k}{t} = 0$). Therefore $Z_{i,t} = \sum_{j\geq i} \binom{t}{j-i} Z_{i,0}$. Since $x_{i,n} = \alpha \cdot p_i \cdot (1 + o(1))$, we have that for every constant $c > 0$, $x_{i,n-c} = \alpha \cdot P_i(c) \cdot (1 + o(1))$ for every $i \geq 2$. In the same way $y_{i,n-c} = \alpha \cdot Q_i(c) \cdot (1 + o(1))$.

Computing $x_{1,t}, y_{1,t}$ (or equivalently $P_{1,t}, N_{1,t}$) needs some care, and this is where several forms of the threshold result are obtained.

Case 1: $\exists j_1, j_2 \geq 2$, $p_{j_1} = n_{j_2} = 1$. The following is the result in this case:

Theorem 4.2 *Let $c > 0$, and let $m = c \cdot n^{k-1}$. Then the probability that* PUR *accepts Φ is equal to $QEMPTY(c \cdot \frac{k!}{\delta_k} \cdot P_2(j))$.*

The proof of the theorem goes along the following lines:

1. as long as $P_{1,t}$ is "small" (sublinear) $P_{1,t-1} \sim P_{1,t} - 1 + Po(t \cdot x_{2,t})$. This is particularly true in the first $\theta(1)$ stages, when $P_{1,t}$ can be approximated by a queue with arrival distribution $Po(c \cdot \frac{k!}{\delta_k} \cdot P_2(n-t))$. This explains the form of the limit probability.
2. Also, in the first $\theta(1)$ stages $P_{1,t}, N_{1,t}$ are "small" (approximately constant), so that w.h.p. PUR does not reject.
3. The probability that PUR accepts after the first $\theta(1)$ stages is small, since, after these stages $P_{1,t}$ will be large enough to make a decrement to 0 unlikely.
4. At the stages $c = n - \theta(\sqrt{n})$, $P_{1,t}, N_{1,t}$ are large enough to guarantee the existence, with nonnegligible probability of a variable that appears both as a positive and a negative unit clause.

Let S be now (Horn $\cup$ 0-valid), $S_H = S \cap HORN$, let Φ be a random formula and Φ_H be its "Horn part". That $SAT(S)$ has the same (coarse) threshold as $SAT(S_H)$ follows easily from the following set of inequalities:

$$\Pr[\Phi \text{ has no positive unit clauses }] \leq$$

$$\Pr[\Phi \in SAT] \leq \Pr[\Phi_H \in SAT].$$

Case 2: $\exists j_1 \geq 2$, $p_{j_1} = 1$ **but** $\forall j \geq 2 : n_j = 0$. Then the following holds:

Theorem 4.3 *Let $c > 0$, and let $m = c \cdot n^{k-1}$. Then the probability that* PUR *accepts Φ is equal to*

$$e^{-c\cdot\frac{k!}{\delta_k}} + (1 - e^{-c\cdot\frac{k!}{\delta_k}}) \cdot QEMPTY(c \cdot \frac{k!}{\delta_k} \cdot P_2(j)).$$

The outline is quite similar to the one of the previous case, with a couple of differences.

1. Now $N_{1,t}$ no longer grows, but remains equal to $N_{1,n}$ for as long as the algorithm does not halt. There exist a nonnegligible (and asymptotically equal to

$e^{-c\cdot\frac{k!}{\delta_k}}$) probability that $N_{1,n} = 0$. In this case $11\ldots 11$ is a satisfying assignment.

2. In the opposite case the structure of the proof (and conclusion) is similar to the one from the Case 1, except that, since $N_{1,t}$ no longer grows, we have to look up to $\theta(n)$ stages to be sure that the algorithm has a nonnegligible probability to reject. In this case the term $\Delta^P_{01,t}$ can no longer be taken to be approximately zero. One can, however, get by, by noticing that, at those stages where $P_{1,t}$ is $\theta(n)$, the probability that there exists a positive unit clause opposite to the negative unit clause guaranteed by the condition $N_{1,n} > 0$ is approximately constant. Iterating this over a small but unbounded number of steps allows us to conclude that for every $\epsilon > 0$ with probability $1 - o(1)$ the formula becomes unsatisfiable in one of the first $\epsilon \cdot n$ stages. Taking ϵ small enough so that $P_{1,t}$ is still nonzero after $\epsilon \cdot n$ stages (if PUR hasn't already stopped by this time) allows us to derive the same form of the limit probability as in case 1.

The analysis of the (Horn $\cup$ 0-valid) case is similar to the previous one.

Case 3: $\exists j_2 \geq 2\ n_{j_2} = 1$ but $\forall j \geq 2 : p_j = 0$.
In this case the threshold result is

Theorem 4.4 *Let $c > 0$, and let $m = c \cdot n^{k-1+\frac{1}{k+1}}$. Then the probability that* PUR *accepts Φ is equal to*

$$e^{-c^{k+1}\cdot(k!)^k} + o(1).$$

The main steps of the analysis are:

1. In this case $P_{1,t}$ is decreasing, but the special form of the threshold makes sure that $\Delta^P_{01,t}$ can be neglected, so $P_{1,t-1} \sim P_{1,t} - 1$, and $P_{1,t} \sim P_{1,n} - (n-t)$.
2. On the other hand $N_{1,t}$ increases and approximately satisfies the following recurrence $N_{1,t-1} \sim N_{1,t} + (t-1)\cdot y_{2,t}$, where $y_{2,t}$ can be computed as outlined before.
3. The probability that the positive literal chosen at stage t occurs both in positive and negative unit form is approximately $1 - e^{-\frac{N_{1,t}}{t}}$.
4. The threshold interval is obtained when the probability that the algorithm rejects in the last $\theta(1)$ stages becomes roughly constant (so that the events "PUR accepts" and "PUR rejects" compete).
5. A recursive computation yields the final form of the limit probability.

An interesting thing happens when considering the corresponding (Horn $\cup$ 0-valid) case: the threshold interval is no longer the one from the corresponding Horn case, but rather mirrors the one in Cases 1 and 2. The underlying reason is simple: the lower bound is the same as in Cases 1 and 2, the probability that Φ contains no positive unit clause. To show an upper bound less than one, consider applying PUR (which is no longer complete) to our formula. With some positive probability PUR will exhaust all the positive unit literals (including those created on the way) before accepting. Since S is not Horn, it contains a clause template with $b \geq 2$ literals.

Such clauses will result, when the positive unit clauses are exhausted, into an at least linear number of clauses of the type $C_{0,b}$. Together with the "all negative" clauses these will ensure that w.h.p. (at least for a big enough constant c) the remaining formula is unsatisfiable. Thus the probability that Φ is satisfiable is less than $1 - \Pr[\text{PUR exhausts all its positive unit clauses}] - o(1)$. The only case left uncovered by this argument is when the only type of "all negative" clauses are the unit clauses, but in this case one can apply a similar reasoning by setting the variables appearing in negative unit clauses too.

c. The argument is based on Friedgut's proof [13] of the fact that 3-SAT has a sharp threshold, and we assume familiarity with the concepts and the methods in this paper. He first shows a general result that roughly states that graph (and hypergraph) problems that have coarse thresholds have a simple approximation at the threshold point. Here is a general and cleaner version of this result from J. Bourgain's appendix:

Proposition 4.5 *Let $A \subset \{0,1\}^n$ be a monotone property, and assume say*

$$\epsilon \leq \mu_p(A) \leq 1-\epsilon$$

$$p\frac{d\mu_p(A)}{dp} < C$$

for some $p = o(1)$ and $C > 0$[2] *Then there is $\delta = \delta(C)$ such that either*

$$\mu_p(\{x \in \{0,1\}^n | x \supset x' \in A, |x'| \leq 10C) > \delta \tag{1}$$

or there exists $x' \notin A$ of size $|x'| \leq 10C$ such that the conditional probability

$$\mu_p(x \in A | x \supset x') > \mu_p(A) + \delta. \tag{2}$$

As a sanity check, let us see how this theorem applies to the three cases of HORN-SAT we have just analyzed. The set A is taken to be $\overline{SAT(S)}$.

- In the first two cases condition 2 applies, and the "magical" formula x' is simply a fixed unit clause.
- In the last case condition 2 applies. The "forbidden formula" x' consists of k different unit clauses $x_1, \ldots, x_k$, together with the clause $\overline{x}_1 \vee \ldots \vee \overline{x}_k$. An unexpected outcome of the analysis is that the satisfiability probability of a random formula Φ *coincides within* $o(1)$ with the probability that Φ contains no isomorphic copy of x'.

[2] such p and C exist, assuming that the sharp threshold condition for A fails with respect to $\epsilon > 0$.

Suppose S is neither (Horn $\cup$ 0-valid) nor (negated Horn $\cup$ 1-valid)

Then S contains the clauses $C_{a_0,0}$ and C_{0,b_0} and $a_0, b_0 \geq 2$. Assume w.l.o.g. that $b_0 \leq a_0$. According to another theorem of Friedgut (that is rederived by Bourgain as Corollary 3), there exists $\gamma \in \mathbf{Q}$ such that the value p from Proposition 4.5 is $\theta(n^{\gamma})$. Therefore the expected number of copies of the clause C_{0,b_0} in a random SAT(S) formula is $\theta(n^{\gamma_1})$, for some rational number γ_1. It is easy to see that $\gamma_1 \geq 0$. Indeed, suppose otherwise. Then the expected number of copies of C_{0,b_0} in Φ is $o(1)$, so with probability $1-o(1)$ Φ contains no clauses consisting of positive literals only. Therefore with probability $1-o(1)$ the assignment 0^n satisfies Φ, which is a contradiction.

Case 1: Suppose $b_{\geq 1} < b_0$.

In this case we want to show that $SAT(S)$ has a sharp threshold. A first observation is that $\gamma_1 > 0$. Indeed, suppose $\gamma_1 > 0$ and consider the formula Ξ obtained from Φ in the following manner: delete from each clause of Φ of length at least b_0 (with probability $1-o(1)$ *all* clauses of Φ are like that) $b_0 - 1$ literals chosen as follows:

- If the clause has at most $b_0 - 1$ positive literals delete them all; then delete a number of random negative literals, so that in the end we delete $b_0 - 1$ literals.
- Otherwise delete all but one of the b_0 positive literals, chosen uniformly at random.

It is easy to see that $\Xi \in SAT \Rightarrow \Phi \in SAT$. Ξ is a Horn formula, falling in the third category (since, by the assumption $b_1 < b_0$ no positive remaining clause has length greater than 1). The formula is *not* a uniform one (since clauses of the same length are *not* do not have the same probability of occurrence). However it can be made so, while increasing the satisfaction probability, by keeping only a fraction of the clauses that occur with probability higher than the minimum one among clauses of the same length. From b. Case 3 it follows that with probability $1-o(1)$ Ξ (therefore Φ) is satisfiable, contradiction.

We are now in position to outline how to mimic Friedgut's argument to show a sharp threshold in our case. Friedgut deals directly with the monotone set A of k-DNF formulas that are tautologies, and first shows that, assuming that this set does not have a sharp threshold it is the alternative 2 that holds. This is evident for K-SAT, but not in our case. Fortunately, we can use some of his argument: assuming that the other alternative holds, the critical value would be $p = \theta(n^{-v/c})$, deriving from an unsatisfiable formula F with v variables and c clauses. To give this threshold, F is also *balanced*, that is, has ratio clauses/variables than any of its induced subformulas. Since F is unsatisfiable it immediately follows that $v < c$. But this cannot happen, since a first moment method easily shows that in our case $p = o(1/n)$.

He then proceeds to show that for k-SAT there cannot exist a "magical" formula x' with the properties guaranteed by Proposition 4.5. The proof follows the following outline (the quotes below refer to statements in [13])

1. the nonexistence of a sharp threshold implies the existence of a small "magical" formula F, which is not itself a tautology, and which boosts the probability that a random formula Φ is a tautology, if we condition on Φ containing a *fixed* copy of F by a non-negligible ($\Omega(1)$) amount.
2. the existence of such a formula implies that adding a constant number of random clauses of size 1 to a random formula also boosts the probability of obtaining a tautology by a positive amount.
3. finally, a contradiction is obtained by showing that were the conclusion of the previous step true, then adding instead an arbitrarily small (but unbounded) number of clauses of size k would also be enough to boost the probability of obtaining a tautology. But such a statement can be refuted directly (Lemma 5.6).

The heart of Friedgut's proof is Step 3, a geometric argument, Lemma 5.7 in his paper. This is where the special syntactical nature of k-SAT (or rather, dually, k-DNF-TAUTOLOGY) appears: according to Lemma 5.7, the probability that an arbitrary subset of the hypercube $\{0,1\}^n$ can be covered with a small (but nonconstant) number of hyperplanes of codimension k (corresponding to DNF-clauses of length exactly k) is asymptotically no smaller than the probability that it can be covered with a constant number of hyperplanes of codimension 1, whose existence is implied by Proposition 4.5 via the process outlined in steps 1,2,3. The clausal structure of $k-SAT$ is reflected by the correspondence between clauses of size k and hyperplanes of codimension k, and this correspondence will extend in our more general case. The argument in Lemma 5.7 is not specific to $k-SAT$, but works in some other cases, if we replace, of course, hyperplanes of codimension k by the corresponding type of hyperplanes and make sure that the geometric argument still works. For instance one can mimic the proof to show that $SAT(S_0)$, where $S_0 = \{C_{a_0,0}, C_{0,b_0}\}$ has a sharp threshold. A minor technical nuisance is that now we need to consider two types of hyperplanes of codimension larger than one, corresponding to both types of clauses, but this does not influence the overall reasoning.

The idea of our argument is now rather transparent: the rest of the steps in Friedgut's argument extend more or less in a straightforward fashion, and it is only the analog of Lemma 5.7 where we need to see how the proof extends. In our case we have a "large" (non-constant) number of copies of $C_{a_0,0}, C_{0,b_0}$ in a random $SAT(S)$ formula at the critical value of p. They are used to "cover a finite number of unit clauses". But this property *does not depend on the other types of clauses in* S, as long as we can make sure that we have a non-constant number of copies of $C_{a_0,0}, C_{0,b_0}$ (this is where $\gamma_1 > 0$ comes into play).

These two types of clauses act as a "$SAT(S_0)$" core of the formula Φ, that is enough to ensure that a the geometric argument used to prove that $SAT(S_0)$ has a sharp threshold holds for

$SAT(S)$ as well. The structure of the proof in this case is similar, at a very high level, with the one of Schaefer's dichotomy theorem: in this latter case the canonical problem is 3-SAT and NP-completeness follows from the ability to "simulate" all clauses of length 3. For sharp/coarse thresholds, the canonical problem is $SAT(S_0)$, and the existence of a sharp threshold follows from the ability to "simulate" both clauses in S_0.

Case 2: Suppose $a_0 = b_0 = \min\{a_{\geq 1}, b_{\geq 1}\}$.

The ideea is similar to the one in Case 1: we show first that the expected number of copies of $C_{a_0,0}$ and C_{0,b_0} is not constant in the critical region, and use Friedgut's argument for S_0. The deletion process is almost identical to the one of the previous section, except that, in order to avoid creating "all negative" clauses of length greater than 1, we do *not* delete the last positive literal, in a clause with less than b_0 positive literals, but a random negative literal.

Case 3.

Assume that we are not into either Case 1 or Case 2 because of the similar inequality for a_0. In this case we want to show that $SAT(S)$ has a coarse threshold, occurring for p such that the expected number of copies of C_{0,b_0} is a constant c. We have already seen that the probability that a random formula Φ is satisfiable is lower bounded by the probabilitythat it contains no copies of C_{0,b_0}. So we only need to argue that the satisfaction probability is strictly less than 1, for some high enough value of the constant in the definition of p.

The main ingredient of this proof, presented in full in the final version of the paper, is the analogue of Friedgut's geometric result to positive clauses only, that allows us to replace w.h.p. a constant number of random unit clauses by a large (but still constant) number of random clauses of length a, while decreasing the satisfaction probability (if at all) by an arbitrarily small constant. Turning this argument "on its head" we infer that for a large enough constant c we can replaces all the clauses counted by C_{0,b_0} by a constant number of unit clauses, while not increasing too much the satisfaction probability.

If $a_0 > b_0$ and $b_{\geq 1} > b_0$, by setting these positive unit clauses to one we create a number of copies of $C_{0,b_{\geq 1}}$ that is at least linear. Together with the existing copies of $C_{a_0,0}$ they insure that (at least for a large enough c) there is a large chance that the remaining formula is unsatisfiable. A similar argument (but working with both positive and negative variables) works for the case $a_0 = b_0 < \min\{a_{\geq 1}, b_{\geq 1}\}$. The other remaining cases are more involved, and rely on an inductive application of the previous idea, but the conclusion is the same, that the satisfaction probability of a random formula is (for large enough c) strictly less than one.

□

5 Conclusions

We have investigated the connection between worst-case complexity and the existence of phase transitions. Our result shows that some connection between the two concepts exists after all: while it is not as clean as the one hoped for in [3], the lack of a phase transition has significant computational implications: such problems are either computationally tractable, or well-predicted by a single, trivial algorithm.

Several open problems remain: a first one is to extend our result to the whole class of generalized satisfiability problems. We believe that obtaining such a characterization is interesting even though the motivating conjecture isn't true. Another question is whether we can extend apply our techniques to constraint programming problems (i.e satisfiability over non-binary domains). Obtaining a complete version of Schaefer's dichotomy theorem in this case is still open; however we believe that some of our results should carry over.

A third, perhaps the most interesting, open question is to elucidate the connection between computational complexity and the "physical" concept of *first-order phase transition*. As we have mentioned, the class of problems with such phase transitions is a subset of the class of problems with sharp thresholds. For clausal generalized satisfiability problems the inclusion is strict: Bollobás et al. [14] have shown that the phase transition in 2-SAT is of second-order. The proof can perhaps be adapted for any (nontrivial) clausal version of 2-SAT, but obtaining any further results is an interesting challenge.

References

[1] G. Istrate. Computational Complexity and Phase Transitions. In preparation. Will be available at http://cnls.lanl.gov/~gistrate/profesional/papers/research.html.

[2] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Determining computational complexity from characteristic phase transitions. *Nature*, 400(8):133–137, 1999.

[3] P. Cheeseman, B. Kanefsky, and W. Taylor. Where the really hard problems are. In *Proceedings of the 11th IJCAI*, pages 331–337, 1991.

[4] T. J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 13th ACM Symposium on Theory of Computing*, pages 216–226. ACM Press, 1978.

[5] Y. Fu. The use and abuse of statistical mechanics in computational complexity. In D. Stein, editor, *Lectures in the Sciences of Complexity*, SFI Series in the Sciences of Complexity. Addison-Wesley Longman, 1989.

[6] S. Mertens. Phase transition in number partitioning problem. *Physical Review Letters*, 81:4281, 1998.

[7] P. Erdős and A. Rényi. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Science*, 5:17–61, 1960.

[8] B. Bollobás and A. Thomason. Threshold functions. *Combinatorica*, 7(1):35–38, 1986.

[9] R. Ladner. On the structure of polynomial time reducibility. *Journal of the ACM*, 22(1):155–171, 1975.

[10] V. Chvátal and B. Reed. Mick gets some (the odds are on his side). In *Proceedings of the 32nd IEEE Symposium on Foundations of Computer Science*, pages 620–626. IEEE Computer Society Press, 1992.

[11] G. Istrate. Critical behavior in the satisfiability of random Horn formulae. manuscript, 1999. A working version is available from http://cnls.lanl.gov/~gistrate/profesional/research.html.

[12] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[13] E. Friedgut. Necessary and sufficient conditions for sharp thresholds of graph properties, and the k-SAT problem. with an appendix by J. Bourgain. *Journal of the A.M.S.*, 12:1017–1054, 1999.

[14] B. Bollobás, C. Borgs, J.T. Chayes, J. H. Kim, and D. B. Wilson. The scaling window of the 2-SAT transition. Technical report, Los Alamos e-print server, http://xxx.lanl.gov/ps/math.CO/9909031, 1999.

Appendix

Proposition 5.1 *For every polynomial time degree $\mathcal{D}$ there exist monotone* NP-decision problems $A, B \in \mathcal{D}$ such that

- A has a coarse threshold.
- B has a sharp threshold.

Proof sketch: Start with two problems $C, D \in P$ that have a coarse (sharp) threshold, for concreteness the property that a graph contains a triangle and 2-UNSAT, respectively). Let $E \in \mathcal{D}$. Encode E into a monotonically increasing set F such that $E \equiv_m^P F$ and $\mu_p(F) \to 1$ as $n \to \infty$ for every p in the "critical region" of C. Define the set A to be the set $C \circ F = \{xy | x \in C, y \in F, |x| = |y|\}$. It is easy too see that $\mu_p(A) = \mu_p(C)(1+o(1))$, so A has a coarse threshold. Moreover $A \in \mathcal{D}$. Set B is constructed in a similar fashion. $\square$