ARTICLES

Rigorous location of phase transitions in hard optimization problems

Dimitris Achlioptas[1], Assaf Naor[1] & Yuval Peres[2]

It is widely believed that for many optimization problems, no algorithm is substantially more efficient than exhaustive search. This means that finding optimal solutions for many practical problems is completely beyond any current or projected computational capacity. To understand the origin of this extreme 'hardness', computer scientists, mathematicians and physicists have been investigating for two decades a connection between computational complexity and phase transitions in random instances of constraint satisfaction problems. Here we present a mathematically rigorous method for locating such phase transitions. Our method works by analysing the distribution of distances between pairs of solutions as constraints are added. By identifying critical behaviour in the evolution of this distribution, we can pinpoint the threshold location for a number of problems, including the two most-studied ones: random *k*-SAT and random graph colouring. Our results prove that the heuristic predictions of statistical physics in this context are essentially correct. Moreover, we establish that random instances of constraint satisfaction problems have solutions well beyond the reach of any analysed algorithm.

Constraint satisfaction problems are at the heart of statistical physics, information theory and computer science. Typically, they involve a large set of variables, each taking values in a small domain, such as {0, 1}, and a collection of constraints, each binding a few of the variables by forbidding some of their possible joint values. Examples include spin-glasses in statistical physics, error-correcting codes in information theory, and satisfiability and graph colouring in computer science.

Given a collection of constraints, a fundamental scientific question is how many of them can be satisfied simultaneously. Value assignments maximizing this number are known as ground states. We are interested in 'polynomial-time' algorithms for finding ground states, that is, algorithms whose running time is bounded by a polynomial in the number of variables. In the *k*-SAT problem there are n binary variables $x_1,\ldots,x_n$, and each constraint (*k*-clause) forbids precisely one out of the 2^k possible values of some $k > 2$ variables; for example, the 3-clause $x_1 \vee x_4 \vee \bar{x}_6$ means that $(x_1, x_4, x_6) \neq (0, 0, 1)$. Trivially, one can determine the ground-states of a *k*-SAT instance in time 2^n, but such an exhaustive search is intractable even when $n = 300$. Many problems of practical interest, such as in chip design, often have $n = 10^5$ or more variables.

Starting with Cook's pioneering work[1], since the 1970s, thousands of problems have been shown to be computationally equivalent to *k*-SAT, from protein-folding to aircraft-crew scheduling. That is, an efficient algorithm for *k*-SAT would immediately give an efficient algorithm for all such problems. It is now widely believed that no efficient algorithm exists for *k*-SAT, that is, that no algorithm can solve all instances efficiently. This is the famous P $\neq$ NP conjecture. At the same time, it is possible that most instances of *k*-SAT can be solved efficiently: perhaps, genuine hardness is only present in a minuscule fraction of all instances. As a result, a major scientific undertaking of the last twenty years has been the study of hardness in typical instances of constraint satisfaction problems (CSPs), generated by sampling uniformly at random among instances with some fixed constraints-to-variables ratio.

A breakthrough[2–5] of the 1990s was the discovery that in typical instances, hardness appears to go along with phase transitions, as suggested in the pioneering work of Fu and Anderson[6] (for more recent accounts see also refs 7 and 8). Specifically, for many CSPs, computational experiments suggest that as constraint density increases, the probability that all constraints can be satisfied drops precipitously from near 1 to near 0; at around the same point, the complexity of finding ground-states appears to increase steeply. In the most-studied example, random instances of *k*-SAT are generated by sampling uniformly, independently and with replacement $m = rn$ constraints from among all possible ones on $x_1,\ldots,x_n$. To understand where the really hard problems are, let us define r_k to be the largest value such that for $r < r_k$, with high probability all constraints can be satisfied. (We will say that an event occurs with high probability (w.h.p.) if its probability tends to 1 as $n \to \infty$.) (Throughout the paper, we will assume that k is arbitrarily large but fixed, while $n \to \infty$.) Similarly, define $r_k^\star$ to be the smallest value such that for $r > r_k^\star$, w.h.p. not all constraints can be satisfied. The Satisfiability Threshold Conjecture asserts that, in fact, $r_k = r_k^\star$ for all $k > 2$ (Fig. 1).

A very simple probabilistic counting argument implies that $r_k^\star \leq 2^k \ln 2$: because constraints are chosen independently, the probability there exists at least one satisfying assignment is at most $2^n (1-2^{-k})^{rn}$, a quantity that tends to 0 for $r \geq 2^k \ln 2$. Heuristic techniques of statistical physics[9–11] also predict that the threshold scales, approximately, as $2^k \ln 2$. On the other hand, all satisfiability algorithms that have been rigorously analysed fail to find satisfying assignments for densities above $c2^k/k$ (we give the bounds corresponding to the best known c in Table 1). This creates a growing chasm between the largest density for which algorithms can provably find solutions and the smallest density for which solutions provably do not exist.

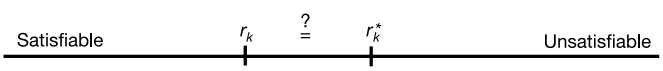

Figure 1 | The Satisfiability Threshold Conjecture. See text for details.

[1]Microsoft Research, One Microsoft Way, Redmond, Washington 98052, USA. [2]Department of Statistics, University of California, Berkeley, California 94720, USA.

Table 1 | Upper and lower bounds for the satisfiability threshold

k	3	4	5	7	10	20	21
Upper bound for $r_k^\star$	4.51	10.23	21.33	87.88	708.94	726,817	1,453,635
Our lower bound for r_k	2.68	7.91	18.79	84.82	704.94	726,809	1,453,626
Algorithmic lower bound for r_k	3.52	5.54	9.63	33.23	172.65	95,263	181,453

The first row gives rigorous[21,22] upper bounds for the location of the satisfiability threshold. The third row gives the largest densities for which some algorithm has been proved to find solutions[23,24]. The second line represents our contribution and gives the largest densities for which we prove that solutions exist. Specifically, we prove that $r_k > 2^k \ln 2 - k$ for all k, and our lower bound converges to $2^k \ln 2 - \frac{k+1}{2} \ln 2 - 1$ as k grows.

Here we resolve this tension by proving that solutions exist much beyond the reach of all analysed algorithms. The key new element of our approach is its focus on how the structure of the space of solutions evolves as density is increased. As a result, we simultaneously get rigorous results on the location of thresholds and insights into why algorithms have such a hard time approaching them. We present the following three concrete results as an illustration of our method.

Statement of results

- We prove that for all $k \geq 3$, the satisfiability threshold lies in the interval $(2^k \ln 2 - k, 2^k \ln 2)$, thus pinpointing its location up to an exponentially small second-order term. Heuristic techniques of statistical physics predict[12] explicit values for the satisfiability threshold for each $k \geq 3$ that scales as $2^k \ln 2 - b_k$, where $b_k \to (1 + \log 2)/2$ (we discuss this point further in the 'Discussion' section). (See Fig. 1.)
- We rigorously determine the asymptotic threshold location for the optimization version of k-SAT, known as 'Max k-SAT', with exponential accuracy. Notably, for this (harder) positive-temperature problem there were no heuristic predictions using the techniques from statistical physics. (See Fig. 2.)

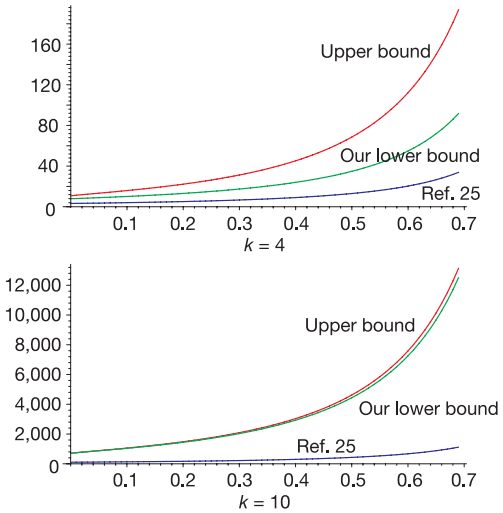

Figure 2 | Our results for random Max k-SAT. Upper and lower bounds for the critical density at which a typical k-SAT instance stops having a truth assignment that satisfies $1 - q2^{-k}$ fraction of its clauses, as a function of $q \in [0, 1)$. The points where the graphs intersect the vertical axis correspond to the bounds for the satisfiability threshold ($q = 0$). The red graph corresponds to the rigorous upper bound $T_k(q) = [2^k \ln 2]/[1 - q + q \ln q]$. The blue graph corresponds to the largest density for which any algorithm has been proved[25] to find assignments satisfying a $1 - q2^{-k}$ fraction of all clauses. Note that this graph is of the order of $T_k(q)/k$, thus rapidly diverging from the red graph with k. In contrast, our lower bound in the green graph converges exponentially fast to the red graph. Specifically, it is of the order of $T_k(q)(1 - \delta_k)$, where $\delta_k = O(k2^{-k/2})$.

- Given a network (graph) G, its chromatic number is the smallest number of colours with which its nodes can be coloured so that adjacent nodes receive different colours. A famous example of graph colouring is the four-colour theorem[13], which states that any planar network (or 'map') has chromatic number at most four. Indeed, in our example below (Fig. 3), three colours suffice. For general, non-planar networks, though, the chromatic number can range anywhere from one to the largest degree plus one.

We prove that this variability is due to a tiny minority of networks. Specifically, we prove that if we pick a graph uniformly at random among all graphs with average degree d, then with probability that tends to 1 as n tends to infinity, its chromatic number is either k_d or $k_d + 1$, where k_d is the smallest integer k such that $d < 2k \ln k$. Thus, fixing the average degree and restricting attention to typical networks can replace planarity in yielding a chromatic number which is essentially known a priori. (See Fig. 3.)

Our method

To resolve whether the failure of algorithms was due to a genuine lack of solutions, as opposed to the difficulty of finding them, we use a method that ignores individual solutions and captures, instead, statistical properties of the entire solution-space. This statistical point of view allows us to avoid the pitfall of computational complexity; we can prove that solutions exist in random instances, without the need to identify a solution for each instance (as algorithms do).

Indeed, if random formulas are genuinely hard near the threshold, then focusing on the existence of solutions rather than their efficient discovery is essential: one cannot expect algorithms to provide accurate results on the threshold's location; they simply cannot get there!

Our approach can be characterized as a 'second moment' method, as it starts from the following basic fact: every non-negative random variable X satisfies $\Pr[X > 0] \geq \mathbb{E}[X]^2/\mathbb{E}[X^2]$. We prove the existence of solutions by applying this inequality to a random variable X that captures an appropriate weighting of the solutions in a CSP. As we will see shortly, such a weighting can be necessary. For example, in random k-SAT, if we simply let X be the number of satisfying assignments, then the ratio $\mathbb{E}[X]^2/\mathbb{E}[X^2]$ is exponentially small in n. An important first step in mitigating this problem was made in ref. 14, where the inequality $r_k \geq 2^{k-1} \ln 2 - O(1)$ was established, by assigning non-zero weight only to those satisfying assignments whose complement is also satisfying.

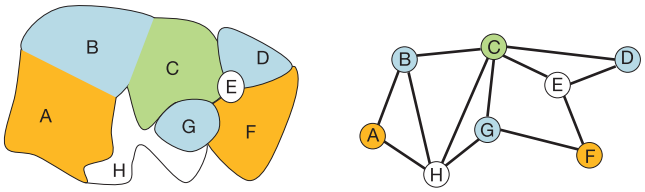

Figure 3 | A planar map and its representation as a network, both properly coloured.

The key to our approach is a systematic search for weights that asymptotically maximize the ratio $\mathbb{E}[X]^2/\mathbb{E}[X^2]$ in the class of tractable weights for which this ratio can be computed. Notably, in choosing these weights we are free to use as hints non-rigorous ideas and insights from physics, without compromising the rigour of the final result. Indeed, while physicists, mathematicians and computer scientists have been investigating the same constraint-satisfaction problems for some time now, this was mostly done using disjoint tool chests. Our approach, on the other hand, can be used in conjunction with the physics heuristics to gain further insight into the geometry of solution spaces; indeed this interaction is already taking place.

The vanilla second moment method fails on random *k*-SAT. Given any k-SAT instance F on n variables, let $X = X(F)$ be its number of satisfying assignments. By computing $\mathbb{E}[X^2]$ and $\mathbb{E}[X]^2$ for random formulas with a given density r, denoted as $F_k(n, rn)$, one can hope to get a lower bound on the probability that $X > 0$, that is, that $F_k(n, rn)$ is satisfiable. Unfortunately, as we show below, this direct application fails dramatically because $\mathbb{E}[X^2]$ is exponentially (in n) greater than $\mathbb{E}[X]^2$ for every density $r > 0$. Nevertheless, this estimation is useful because it points to the source of the problem and lays the foundation for our later successful choice of X.

For a k-CNF formula with independent clauses $c_1, c_2, \ldots, c_m$ it is straightforward to show that:

$$\mathbb{E}[X^2] = 2^n \sum_{z=0}^{n} \binom{n}{z} f_S(z/n)^m \qquad (1)$$

where $f_S(\alpha) = 1 - 2^{1-k} + 2^{-k}\alpha^k$ is the probability that two fixed truth assignments that agree on $z = \alpha n$ variables both satisfy a randomly drawn clause. Observe that f is an increasing function of α and that $f_S(1/2) = (1 - 2^{-k})^2$, which means that truth assignments having overlap $n/2$ are uncorrelated.

Letting $\Lambda_S(\alpha) = [2f_S(\alpha)^r]/[\alpha^\alpha(1-\alpha)^{1-\alpha}]$ we see that $\mathbb{E}[X]^2 = (2^n(1-2^{-k})^{rn})^2 = (4f_S(1/2)^r)^n = \Lambda_S(1/2)^n$. Because $\binom{n}{\alpha n} = (\alpha^\alpha(1-\alpha)^{1-\alpha})^{-n} \times \mathrm{poly}(n)$ we see from equation (1) that $\mathbb{E}[X^2] \geq (\max_{0\leq\alpha\leq1}\Lambda_S(\alpha))^n \times \mathrm{poly}(n)$. Therefore, if there exists some $\alpha \neq 1/2$ such that $\Lambda_S(\alpha) > \Lambda_S(1/2)$, then the second moment is exponentially greater than the square of the expectation and we get only an exponentially small lower bound for $\Pr[X > 0]$. Put differently, unless the dominant contribution to $\mathbb{E}[X^2]$ comes from uncorrelated pairs of satisfying assignments, that is, pairs with overlap $n/2$, the second moment method fails. Unfortunately, for all $r > 0$, we have $\Lambda_S'(1/2) \neq 0$. This is because the entropic factor $\varepsilon(\alpha) = 1/(\alpha^\alpha(1-\alpha)^{1-\alpha})$ is symmetric around $\alpha = 1/2$, while f_S is increasing in $(0, 1)$. As a result, the derivative of Λ_S becomes 0 only when the correlation benefit balances with the penalty of decreasing entropy at some $\alpha > 1/2$. We demonstrated this, for $k = 5$, in Fig. 4.

In general, given a real number $\alpha \in [0, 1]$, we would like to know the number of pairs of satisfying truth assignments that agree on $z = \alpha n$ variables in a random formula. Each term in the sum in equation (1) gives us the expected number of such pairs. Although this expectation overemphasizes formulas with more satisfying assignments, it gives valuable information on the distribution of distances among truth assignments in a random formula. For example, if for some values of z (and k, r) this expectation tends to 0 with n, we can infer that w.h.p. there are no pairs of truth assignments that agree on z variables in a random k-CNF formula with density r.

Balance and the weighted second moment method. An attractive feature of the second moment method is that we are free to apply it to any random variable $X = X(F)$ such that $X > 0$ implies F is satisfiable. Sums of the form $X = \Sigma_\sigma w(\sigma, F)$ clearly have this property as long as $w(\sigma, F) = 0$ when σ is not a satisfying assignment. Such weighting schemes can be viewed as transforms of the original problem and can be particularly effective in exploiting insights into the source of correlations.

With this in mind, let us consider random variables of the form $X = \Sigma_\sigma \Pi_c w(\sigma, c)$, where w is some arbitrary function. (Eventually, we will require that $w(\sigma, c) = 0$ if σ falsifies c.) Similarly to equation (1), it is rather straightforward to prove that $\mathbb{E}[X^2] = 2^n \Sigma_{z=0}^{n} \binom{n}{z} f_w(z/n)^m$, where $f_w(z/n) = \mathbb{E}[w(\sigma, c)\, w(\tau, c)]$ is the correlation between two truth assignments σ and τ that agree on z variables, with respect to a single random clause c. It is also not hard to see that $f_w(1/2) = \mathbb{E}[w(\sigma, c)]^2$, that is, truth assignments at distance $n/2$ are uncorrelated for any function w. Thus, arguing as in the previous section, we see that $\mathbb{E}[X^2]$ is exponentially greater than $\mathbb{E}[X]^2$ unless $f_w'(1/2) = 0$.

At this point we observe that because we are interested in random formulas where literals are drawn uniformly, it suffices to consider functions w such that: for every truth assignment σ and every clause $c = \ell_1 \vee \cdots \vee \ell_k$, $w(\sigma, c) = w(\mathbf{v})$, where $v_i = +1$ if ℓ_i is satisfied under σ and $v_i = -1$ if ℓ_i is not satisfied under σ. (So, we will require that $w(-1, \ldots, -1) = 0$.) Letting $A = \{-1, +1\}^k$ and differentiating f_w yields the geometric condition:

$$f_w'(1/2) = 0 \Leftrightarrow \sum_{\mathbf{v}\in A} w(\mathbf{v})\mathbf{v} = 0 \qquad (2)$$

The condition in the right-hand side (r.h.s.) of equation (2) asserts that the vectors in A, when scaled by $w(\mathbf{v})$, must cancel. This gives us another perspective on the failure of the vanilla second moment method: when $w = w_S$ is the indicator variable for satisfiability, the condition in the r.h.s. of equation (2) does not hold because the

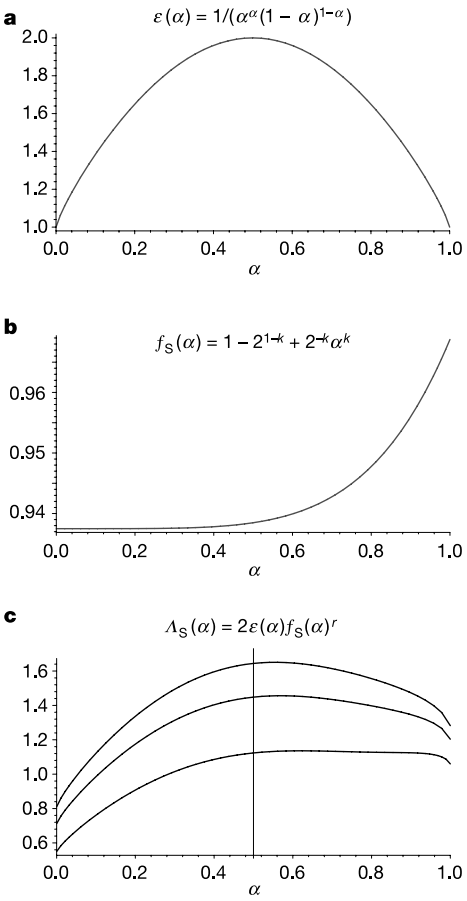

Figure 4 | Plots of entropy, correlation, and their product for the vanilla second moment method. The function f_S is plotted for $k = 5$. The function Λ_S is plotted for $k = 5$ with different values of r. Specifically, $r = 14, 16, 20$ from top to bottom. The vertical line at $\alpha = 1/2$ highlights that Λ_S is maximized for $\alpha > 1/2$.

vector $(-1, -1, \ldots, -1)$ has weight 0, while all other $\mathbf{v} \in A$ have weight 1.

Note that the r.h.s. of equation (2) implies that in a successful weighting each coordinate must have mean 0: each literal must be equally likely to be $+1$ or -1 when we pick truth assignments with probability proportional to their weight under w. We will call truth assignments with $km/2$ satisfied literal occurrences 'balanced'.

To make the second moment method work we would like to choose a function w that is 'as close as possible' to w_S while being balanced, that is, satisfy equation (2). For $\mathbf{v} \in A$, let $|\mathbf{v}|$ denote its number of $+1$ values. Maximizing the relative entropy of w with respect to w_S subject to equation (2) yields:

$$w(\mathbf{v}) = \begin{cases} 0 & \text{if } \mathbf{v} = (-1, \ldots, -1) \\ \lambda^{|\mathbf{v}|} & \text{otherwise} \end{cases} \quad (3)$$

where λ satisfies $(1+\lambda)^{k-1} = 1/(1-\lambda)$. In Fig. 5 we plot the functions f_w and Λ_w corresponding to the w in equation (3) for the same choice of k, r as in Fig. 4. For $k \geq 3$, we prove[15] that Λ_w is maximized at $\alpha = 1/2$ as long as $r \leq 2^k \ln 2 - (k+1)\frac{\ln 2}{2} - 1 - \beta_k$, where $\beta_k \to 0$. This implies that for such r, $\mathbb{E}[X^2] < C \cdot \mathbb{E}[X]^2$, where $C = C(k) > 0$ is independent of n. By the second moment method this implies $\Pr[X > 0] \geq 1/C$ and by a result of Friedgut[16], it follows that $r_k \geq r$.

To gain additional insight into balanced assignments it helps to think of $F_k(n, m)$ as generated in two steps: first choosing the km literal occurrences randomly, and then partitioning them into clauses. Already, at the end of the first step, truth assignments that satisfy many literal occurrences have significantly greater conditional probability of being satisfying. But such assignments are highly correlated because in order to satisfy many literal occurrences they tend to agree with the majority truth assignment, and thus each other, on more than half the variables. Our choice of λ penalizes satisfying assignments that satisfy more than half of all literal occurrences in the formula: that is, more than a random assignment. In other words, our random variable X curbs the tendency of satisfying assignments to lean towards the majority vote assignment. The fact that Λ_w is maximized at $\alpha = 1/2$ for densities almost all the way to the random k-SAT threshold, means that for all such densities the expected number of pairs of balanced assignments at distance $n/2$ is exponentially greater than the expected number of pairs of balanced assignments at any other distance.

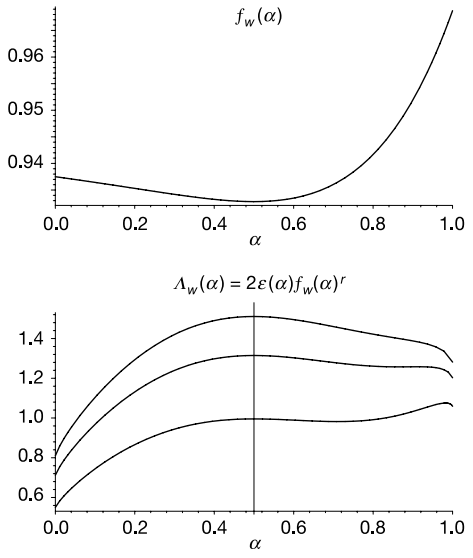

Figure 5 | Plots of the correlation function and its product with entropy for the weighted second moment method, when the weighing is given by equation (3). The function f_w is plotted for $k = 5$. The function Λ_w is plotted for $k = 5$ with the same values of r as in Fig. 2, namely, $r = 14, 16, 20$ from top to bottom. The vertical line at $\alpha = 1/2$ shows that Λ_w is maximized at $\alpha = 1/2$ for $r = 14$ and $r = 16$, but at some $\alpha > 1/2$ for $r = 20$.

Random Max k-SAT

Say that a k-SAT instance with m clauses is 'p-satisfiable', where $p \in [0, 1]$, if there exists an assignment satisfying at least $(1 - 2^{-k} + p2^{-k})m$ clauses (observe that every k-CNF formula is 0-satisfiable since a random truth assignment satisfies $(1 - 2^{-k})m$ clauses in expectation). We define $r_k(p)$ to be the largest value such that for $r < r_k(p)$, w.h.p. a random formula $F_k(n, rn)$ is p-satisfiable. Similarly, we define $r_k^*(p)$ to be the smallest value such that for $r > r_k^*(p)$, w.h.p. a random formula $F_k(n, rn)$ is not p-satisfiable. So, in these terms, the Satisfiability Threshold Conjecture amounts to $r_k(1) = r_k^*(1)$.

As we saw earlier for the case $p = 1$, that is, for satisfiability, a major factor in the excessive correlations behind the failure of the vanilla second moment method is that satisfying truth assignments tend to lean toward the majority vote truth assignment. To avoid this pitfall, again, we consider balanced truth assignments, this time p-satisfying ones. Specifically, for $\sigma \in \{0, 1\}^n$ we let:

(1) $H = H(\sigma, F)$ be the number of satisfied literal occurrences in F under σ, minus the number of unsatisfied literal occurrences in F under σ, and

(2) $U = U(\sigma, F)$ be the number of unsatisfied clauses in F under σ, minus $m(1-p)/2^k$.

To focus on the desired truth assignments we fix $\beta > 0$ and $0 < \gamma < 1$ and define $X(\beta, \gamma)$ as:

$$X(\beta, \gamma) = \sum_{\sigma} \gamma^{H(\sigma,F)} e^{-\beta U(\sigma,F)} \quad (4)$$

Because $\beta > 0$ and $0 < \gamma < 1$ we see that the truth assignments σ for which $H(\sigma, F) > 0$ or $U(\sigma, F) > 0$ are suppressed exponentially, while the rest are rewarded exponentially. In statistical physics terms, β can be interpreted as an inverse temperature, where perfect satisfiability, analysed in the previous section, corresponds to 0 temperature, that is, $\beta \to \infty$.

By applying the second moment method to X we pinpoint[17] the values of $r_k(p)$ and $r_k^*(p)$ with relative error that tends to zero exponentially fast in k. Specifically, for every $p \in (0, 1)$ let:

$$T_k(p) = \frac{2^k \ln 2}{p + (1-p)\ln(1-p)}$$

and define $T_k(1) = 2^k \ln 2$ so that $T_k(\cdot)$ is continuous on $(0, 1]$.

Theorem 1

There exists a sequence $\delta_k = O(k2^{-k/2})$, such that for all $k \geq 2$ and $p \in (0, 1]$:

$$(1 - \delta_k)T_k(p) < r_k(p) \leq r_k^*(p) < T_k(p)$$

Random graph colouring

To generate a typical network with n nodes and average degree d, we start with n isolated nodes and join each pair of them with probability $p = d/(n-1)$, independently of all others. This is known as the Erdös–Rényi $G(n, p)$ model of random graphs. We prove:

Theorem 2

For any real number $d > 0$, let k_d be the smallest integer k such that $d < 2k \ln k$. With high probability the chromatic number of a random $G(n, p)$ graph with average degree d is either k_d or $k_d + 1$.

To prove Theorem 2 we start with n isolated nodes and repeat m times: pick two nodes uniformly, independently, with replacement, and join them. We claim that for every $d > 0$, if $m = dn/2$, then w.h.p. the chromatic number of the resulting (multi)graph G is either

k_d or $k_d + 1$. By standard results of random graph theory[18], this implies Theorem 2. To prove our claim we first observe that the probability that G is k-colourable is at most $k^n(1 - 1/k)^{dn/2}$, which tends to 0 for $d > 2k \ln k$. Our main contribution is to prove that for slightly smaller d, namely $d < 2(k - 1) \ln(k - 1)$, w.h.p. G is k-colourable. To do that we proceed as follows.

Let $X = X(G)$ be the number of balanced k-colourings of G, that is, k-colourings in which all colour classes have equal size. We apply the second moment method to X. To compute $\mathbb{E}[X^2]$ we consider all pairs σ, τ of candidate solutions (all pairs of balanced k-partitions of the n vertices) and for each such pair determine the probability that both k-partitions will assign distinct colours to the endpoints of a randomly drawn edge. If $a_{ij}n$ is the number of vertices having colour i in σ and colour j in τ it is not hard to see that this probability is:

$$1 - \frac{2}{k} + \sum_{i,j} a_{ij}^2 \tag{5}$$

Equation (5) highlights the main difficulty we need to overcome in this context: whereas in satisfiability problems the overlap parameter for a pair of solutions was a single integer (their Hamming distance), now it is a $k \times k$ matrix. As a result, to determine $\mathbb{E}[X^2]$ we need to resolve an entropy–energy tradeoff over doubly stochastic matrices, a problem of major analytic difficulty.

More precisely, we show that $\mathbb{E}[X^2]$ is dominated by the contribution of the pairs of k-partitions whose overlap matrix A maximizes the function:

$$g_d(A) = -\sum_{i,j} a_{ij} \log a_{ij} + \frac{d}{2} \cdot \log\left(1 - \frac{2}{k} + \sum_{i,j} a_{ij}^2\right)$$

The first term in g_d measures the number of pairs of k-partitions that have A as their overlap matrix, while the second term measures the probability that such pairs will be valid k-colourings in a random graph with average degree d. When $a_{ij} = 1/k^2$ for all i, j, corresponding to uncorrelated partitions, the first term is maximized while the second is minimized. At the other extreme, when each row and column has precisely one element equal to $1/k$, corresponding to perfectly correlated partitions, the first term is minimized while the second term is maximized. If we interpret A as spreading a fixed amount of mass over k^2 cells, we see that there are two 'forces' that determine its shape: entropy, favouring flatness, and energy, favouring the formation of peaks.

To resolve this entropy–energy tradeoff we develop general optimization tools that we expect to be of much broader applicability. In particular, we prove that as d is increased the maximizer switches instantaneously from the perfectly flat matrix J_k to a matrix in which more than half the mass is captured by only k entries. At that point the dominant contribution to $\mathbb{E}[X^2]$ stops corresponding to uncorrelated k-colourings and the second moment fails. Our result follows by proving[19] that the switch in the locus of the maximizer occurs for some $d > 2(k - 1) \ln(k - 1)$.

Discussion

Our work implies that the physics predictions for certain key combinatorial optimization and decision problems can be rigorously justified. Indeed, recently, non-rigorous techniques of statistical physics were used to derive[11] predictions for the location of the threshold for random k-SAT for every fixed $k \geq 3$; for example, for $k = 3$, 4, 5 the predicted values are 4.267, 9.931 and 21.117, respectively. These predictions are compatible with the rigorous bounds and indeed, match the rigorous upper bound as $k \to \infty$. Similarly, for the graph colouring problem it was predicted in ref. 20 that the k-colourability transition occurs at $d = 2k \log k - 1 + o(1)$, which fits neatly between the rigorous upper and lower bounds we state in the 'Random graph colouring' section.

The gap between our lower bound for the location of the k-SAT threshold and the best-known algorithmic lower bound $r_k = \Omega(2^k/k)$, seems to us the most significant remaining problem. Indeed, the physics ideas mentioned above have motivated a new satisfiability algorithm[10,11] that performs extremely well for small values of k, such as $k = 3$. A rigorous analysis of this algorithm is still lacking, and it remains unclear whether its success for values of r close to the threshold extends to large k. Indeed, even evaluating the algorithm experimentally is already intractable for moderate values of k, such as $k = 10$.

The success of the second moment method for balanced satisfying truth assignments suggests that such assignments form a 'mist' in $\{0, 1\}^n$ and, as a result, they might be hard to find by algorithms based on local updates. (More precisely, the satisfying assignments decompose into clusters whose diameter is much smaller than the inter-cluster distances, as predicted in ref. 11.) Moreover, as k increases the influence exerted by the majority vote assignment becomes less and less significant as most literals occur very close to their expected $kr/2$ times. As a result, as k increases, typical satisfying assignments get closer and closer to being balanced, meaning that the structure of the space of solutions for small values of k might be significantly different from the structure for large values of k. Indeed, an appealing intuitive explanation for the fact that our methods succeed where all analysed algorithms fail is that the latter rely on knowing part of a solution to get an indication on the values of the unknown variables. The success of such a strategy inevitably requires that the space of solutions is clustered, which means that solutions are highly correlated. Our method, on the other hand, targets situations in which solutions form a 'sparse mist' in the configuration space. Thus, it draws strength precisely from the phenomenon that causes algorithms to fail.

To summarize, the following key question remains: is there an algorithmic threshold $\lambda_k = o(2^k)$ so that for densities $r > \lambda_k$, no polynomial-time algorithm can find a satisfying truth assignment with probability bounded away from 0 as $n \to \infty$?

Received 16 December 2004; accepted 31 March 2005.

1. Cook, S. A. The complexity of theorem-proving procedures. *Proc. 3rd Ann. ACM Symp. on Theory of Computing* 151–158 (1971).
2. Cheeseman, P., Kanefsky, B. & Taylor, W. Where the really hard problems are. *Proc. 12th Int. Joint Conf. on Artificial Intelligence* 331–337 (1991).
3. Selman, B., Levesque, H. & Mitchell, D. Hard and easy distributions of SAT problems. *Proc. 10th Nat. Conf. on Artificial Intelligence* 459–465 (1992).
4. Kirkpatrick, S. & Selman, B. Critical behavior in the satisfiability of random boolean expressions. *Science* **264,** 1297–1301 (1994).
5. Monasson, R., Zecchina, R., Kirkpatrick, S., Selman, B. & Troyansky, L. Determining computational complexity from characteristic 'phase transitions'. *Nature* **400,** 133–137 (1999).
6. Fu, Y. & Anderson, P. W. Application of statistical mechanics to NP-complete problems in combinatorial optimisation. *J. Phys. A* **19,** 1605–1620 (1986).
7. Anderson, P. W. Solving problems in finite time. *Nature* **400,** 115–116 (1999).
8. Gomes, C. P. & Selman, B. Satisfied with physics. *Science* **297,** 784–785 (2002).
9. Monasson, R. & Zecchina, R. Statistical mechanics of the random K-satisfiability model. *Phys. Rev. E* **56,** 1357–1370 (1997).
10. Mézard, M. & Zecchina, R. Random K-satisfiability: from an analytic solution to a new efficient algorithm. *Phys. Rev. E* **66,** 056126 (2002).
11. Mézard, M., Parisi, G. & Zecchina, R. Analytic and algorithmic solution of random satisfiability problems. *Science* **297,** 812–815 (2002).
12. Mertens, S., Mézard, M. & Zecchina, R. Threshold values of random k-SAT from the cavity method. To appear in *Random Struct. Algorithms* (in the press); preprint at http://arxiv.org/abs/cs.CC/0309020 (2003).
13. Appel, K. & Haken, W. Every planar map is four colorable. *Contemp. Math.* **98,** (1989).
14. Achlioptas, D. & Moore, C. The asymptotic order of the random k-SAT threshold. *Proc. 43rd Ann. IEEE Symp. on Foundations of Computer Science* 779–788 (2002).
15. Achlioptas, D. & Peres, Y. The threshold for random k-SAT is $2^k \log 2 - O(k)$. *J. Am. Math. Soc.* **17,** 947–973 (2004).
16. Friedgut, E. Sharp thresholds of graph properties, and the k-SAT problem. *J. Am. Math. Soc.* **12,** 1017–1054 (1999).
17. Achlioptas, D., Naor, A. & Peres, Y. On the maximum satisfiability of random formulas. *Proc. 44th Ann. IEEE Symp. on Foundations of Computer Science* 362–370 (2003).

18. Bollobás, B. *Random Graphs* Vol. 73, 2nd edn *Cambridge Studies in Advanced Mathematics* (Cambridge Univ. Press, Cambridge, 2001).
19. Achlioptas, D. & Naor, A. The two possible values of the chromatic number of a random graph. *Proc. 36th Ann. ACM Symp. on Theory of Computing* 587–593 (2004).
20. Krzakala, F., Pagnani, A. & Weigt, M. Threshold values, stability analysis and high-q asymptotics for the coloring problem on random graphs. *Phys. Rev. E* **70,** 046705 (2004).
21. Kirousis, L. M., Kranakis, E., Krizanc, D. & Stamatiou, Y. Approximating the unsatisfiability threshold of random formulas. *Random Struct. Algorithms* **12,** 253–269 (1998).
22. Dubois, O., Boufkhad, Y. & Mandler, J. Typical random 3-SAT formulae and the satisfiability threshold. *Proc. 11th Ann. ACM-SIAM Symp. on Discrete Algorithms* 126–127 (2000).
23. Kaporis, A. C., Kirousis, L. M. & Lalas, E. G. The probabilistic analysis of a greedy satisfiability algorithm. *Proc. 10th Ann. European Symp. on Algorithms* 574–585 (2002).
24. Frieze, A. M. & Suen, S. Analysis of two simple heuristics on a random instance of *k*-SAT. *J. Algorithms* **20,** 312–355 (1996).
25. Coppersmith, D., Gamarnik, D., Hajiaghayi, M. T. & Sorkin, G. B. Random MAX SAT random MAX CUT, and their phase transitions. *Random Struct. Algorithms* **24,** 502–545 (2004).

Acknowledgements We thank S. Kirkpatrick, J. Kleinberg and S. Mertens for feedback on the presentation of the results.

Author Information Reprints and permissions information is available at npg.nature.com/reprintsandpermissions. The authors declare no competing financial interests. Correspondence and requests for materials should be addressed to D.A. (optas@microsoft.com).