

# Determining computational complexity from characteristic 'phase transitions'

Rémi Monasson\*, Riccardo Zecchina†, Scott Kirkpatrick‡, Bart Selman§ & Lidror Troyansky||

\* CNRS-Laboratoire de Physique Théorique, 24, Rue Lhomond, 75231 Paris Cedex, France

† The Abdus Salam International Centre for Theoretical Physics, Strada Costieri 11, 34100 Trieste, Italy

‡ IBM, Thomas J. Watson Research Center, Yorktown Heights, New York 10598, USA

§ Computer Science Department, Cornell University, Ithaca, New York 14853, USA

|| Institute of Computer Science, The Hebrew University, Jerusalem 91904, Israel

**Non-deterministic polynomial time (commonly termed 'NP-complete') problems are relevant to many computational tasks of practical interest—such as the 'travelling salesman problem'—but are difficult to solve: the computing time grows exponentially with problem size in the worst case. It has recently been shown that these problems exhibit 'phase boundaries', across which dramatic changes occur in the computational difficulty and solution character—the problems become easier to solve away from the boundary. Here we report an analytic solution and experimental investigation of the phase transition in  $K$ -satisfiability, an archetypal NP-complete problem. Depending on the input parameters, the computing time may grow exponentially or polynomially with problem size; in the former case, we observe a discontinuous transition, whereas in the latter case a continuous (second-order) transition is found. The nature of these transitions may explain the differing computational costs, and suggests directions for improving the efficiency of search algorithms. Similar types of transition should occur in other combinatorial problems and in glassy or granular materials, thereby strengthening the link between computational models and properties of physical systems.**

Many computational tasks of practical interest are surprisingly difficult to solve even using the fastest available machines. Such problems, found for example in planning, scheduling, machine learning, hardware design, and computational biology, generally belong to the class of NP-complete problems<sup>1–3</sup>. NP stands for 'non-deterministic polynomial time', which denotes an abstract computational model with a rather technical definition. Intuitively speaking, this class of computational tasks consists of problems for which a potential solution can be checked efficiently for correctness, yet finding such a solution appears to require exponential time in the worst case. A good analogy can be drawn from mathematics: proving open conjectures in mathematics is extremely difficult, but verifying any given proof (or solution) is generally relatively straightforward.

The class of NP-complete problems lies at the foundations of the theory of computational complexity in modern computer science. Literally thousands of computational problems have been shown to be NP-complete. The completeness property of NP-complete problems means that if an efficient algorithm for solving just one of these problems could be found, one would immediately have an efficient algorithm for all NP-complete problems. However, a fundamental conjecture of modern complexity theory is that no such efficient algorithm exists.


Although NP-complete problems are believed to require exponential time to solve in the worst case, the typical-case behaviour is difficult to characterize. Yet, such typical-case properties are most relevant in practical applications. Fu and Anderson<sup>4</sup> first conjectured a deep connection between NP-complete problems and models studied in statistical physics<sup>5</sup>. More recently, it was discovered (refs 6–9) that NP-complete problems can exhibit phase transition phenomena, analogous to those in physical systems, with the hardest problems occurring at the phase boundary. (For a collection of papers and a review, see refs 10, 11.) The exact relationship between phase transition phenomena and computational properties has remained unclear. For example, phase

transitions are also observed in computationally 'easy' problems (that is, ones that are not NP-complete).

Here we establish the nature of the relationship between phase transition phenomena and typical-case computational complexity. More specifically, we show that when the underlying computational task exhibits a continuous ('second order') phase transition, resource requirements grow only polynomially with problem size, while discontinuous ('random first order'<sup>12</sup>) phase transitions correspond to an exponential growth in resource requirements, characteristic of truly hard computational problems. Our results show that techniques from statistical physics can provide important new insights into computational phenomena, possibly leading to improved solution methods. In addition, computational problems, viewed as many-particle systems, provide models which challenge the assumptions of physics, and may shed light on the behaviour of some highly nonlinear materials now coming under study.

## **$K$ -SAT: a standard test bed**

We focus on the  $K$ -satisfiability ( $K$ -SAT) problem<sup>1–3,11</sup>, which is characteristic of the class of NP-complete problems, and is commonly used as a test bed for heuristic algorithms. An instance of the  $K$ -SAT problem is defined by a set of  $N$  boolean variables, each of which can be assigned 'true' or 'false', and a set of  $M$  clauses. Each clause is a logical OR of  $K$  variables, where each variable can be negated. The goal is to find an assignment to the variables such that all clauses are satisfied. For example, the 2-SAT formula,  $(x \text{ OR } y) \text{ AND } ((\text{NOT } x) \text{ OR } (\text{NOT } y))$ , can be satisfied by setting  $x$  to true and  $y$  to false. In the worst case, for  $K \geq 3$ , an exponential search through the space of  $2^N$  truth assignments is required in order to find a satisfying assignment or to determine that no such assignment exists, in which case the formula is called 'unsatisfiable'. The  $K$ -SAT problem was the first problem shown to be NP-complete<sup>1</sup>. A remarkable consequence of the theory of NP-completeness is that a large collection of problems (in fact, several thousand) from many different fields can be represented as instances of the  $K$ -SAT



**Figure 1** The SAT/UNSAT phase transition for the  $(2 + p)$  SAT problems for a range of values of  $p$  and  $N$ . Each curve was generated by considering 10,000–40,000 randomly generated formulae, for a range of values of  $\alpha$ . For  $p = 0$ ,  $N$  has values of 1,000 (red), 2,000 (blue), 5,000 (green), 10,000 (light blue); for  $p = 0.2$ ,  $N$  has values of 1,000 (red), 2,000 (blue), 5,000 (green), 7,500 (light blue); for  $p = 0.4$ ,  $N$  has values of 500 (red), 1,000 (blue), 2,000 (green), 5,000 (light blue); for  $p = 0.6$ ,  $N$  has values of 250 (red), 500 (blue), 1,000 (green), 1,500 (light blue);  $p = 0.8$ ,  $N$  has

values of 200 (red), 350 (blue), 500 (green); for  $p = 1.0$ ,  $N$  has values of 100 (red), 200 (blue), 250 (green). The largest value of  $N$  which could be achieved for each value of  $p$  was limited by the increasing cost of the calculations as  $p$  is increased. We note that  $p = 0.0$  corresponds to 2-SAT, and  $p = 1.0$  to 3-SAT. The vertical lines mark the thresholds as predicted by the replica symmetric (RS) theory. The transition curves sharpen with increasing  $N$ .

problem<sup>2</sup>. In a sense, the problem lies at the core of computational complexity theory, and the analysis of  $K$ -SAT therefore provides insights into a broad range of challenging computational problems.

We study the typical case complexity of  $K$ -SAT by considering an ensemble of randomly generated  $K$ -SAT formulae. For each formula, we generate  $M$  clauses, where each clause is generated by randomly selecting  $K$  variables from the set of  $N$  variables and negating each selected variable with probability 0.5 to construct the clause. Formulae constructed at random, keeping the ratio of clauses to variables,  $\alpha = M/N$ , constant as  $M, N \rightarrow \infty$ , provide a natural ensemble of test problems.

The value of  $K$  is important: the 1-SAT and 2-SAT problem can be solved efficiently, that is, in time polynomial in size of the formula (in fact, a linear time algorithm is known<sup>13</sup>). We say that 1-SAT and 2-SAT belong to the class ‘P’ (polynomial time solvable problems). The 2-SAT problem exhibits phase transition phenomenon at a critical ratio of clauses to variables equal to one ( $\alpha_c(2) = 1$ ). Below this critical ratio almost all formulae are satisfiable (SAT), whereas above this ratio almost all are unsatisfiable (UNSAT)<sup>14,15</sup>. The intuitive explanation for this phenomenon is that at low values of  $\alpha$ , we have relatively few clauses and many variables, which means that the formula is relatively easy to satisfy (the problem is under-constrained). At high values of  $\alpha$ , the problem is over-constrained and generally no satisfying assignment exists. The most interesting formulae can be found at the transition from the under-constrained to the over-constrained area. We will make these observations more precise below.

For  $K \geq 3$ ,  $K$ -SAT is NP-complete. Using computer experiments, we have shown<sup>8</sup> thresholds for  $K = 3, 4, 5$  and 6. A threshold exists for any value of  $K$  (ref. 16), but determining the exact location of the threshold remains a difficult open problem. One reason for the interest in the threshold is the growing recognition that ‘easy–hard–easy’ patterns of computational cost are characteristic of heuristic search algorithms developed for solving NP-complete problems,

and that the hardest instances occur near phase boundaries<sup>6–11</sup>. (The easy–hard–easy pattern refers to the facts that (1) at low values of  $\alpha$ , it is relatively easy to find a satisfying assignment, (2) at values of  $\alpha$  near the threshold, it is most difficult to find a satisfying assignment or show unsatisfiability of the formula, and (3) at high values of  $\alpha$ , it becomes again relatively easy to solve the formulae (in this case, by showing them unsatisfiable).) Above the critical ratio,  $\alpha_c$ , proving UNSAT is still hard, with a lower bound on the search cost that is exponential in  $N$  with a coefficient of  $N$  that decreases as a power law in  $\alpha$  (ref. 17). Because of the extreme hardness of formulae at the threshold, such formulae provide a useful benchmark for evaluating new search methods.

### Interpolating between P and NP

To understand the onset of exponential complexity that occurs when going from a problem in P (2-SAT) to a problem that is NP-complete (3-SAT), we introduce here formulae containing mixtures of 2- and 3-clauses. Consider a random formula with  $M$  clauses, of which  $(1 - p)M$  contain two variables and  $pM$  contain three variables, with  $0 \leq p \leq 1$ . This ‘ $(2 + p)$ -SAT’ model smoothly interpolates between 2-SAT ( $p = 0$ ) and 3-SAT ( $p = 1$ ). Figure 1 shows that the threshold from SAT to UNSAT shifts as a function of  $p$  from that for 2-SAT to that for 3-SAT. The problem is NP-complete for any value of  $p > 0$ , as any such instance contains a sub-formula of 3-clauses. Thus, the worst-case complexity is probably exponential, but the typical-case complexity need not be, as we find below.


There is a strong analogy between randomly generated  $K$ -SAT problems and disordered materials, alloys or glasses, studied by constructing models whose energy function captures the unsatisfied constraints. In studies of strongly disordered models with conflicting interactions, similar to the randomly negated variables in  $K$ -SAT, the replica method is used<sup>5</sup> to confirm that ordering of an unusual type is possible in the presence of microscopic randomness. We have previously applied replica methods to determine the

characteristics of the 3-SAT problem<sup>18–20</sup>. Here we report results on the  $(2 + p)$ -SAT problem, which provides new insights into the typical-case complexity, in particular with respect to the change in going from a computationally tractable (polynomial) problem to an intractable (NP-complete) problem. We hope that these results are useful both to computer scientists, in constructing better algorithms, and to statistical physicists, in their study of disordered systems.

The  $(2 + p)$ -SAT model can be mapped onto a diluted spin glass model with  $N$  spin variables  $S_i$ :  $S_i = 1$  if the boolean variable  $x_i$  is true,  $S_i = -1$  if  $x_i$  is false. We can associate with each configuration an energy  $E$ , or cost-function, equal to the number of clauses violated by that configuration. (Note that a configuration corresponds to an assignment of truth values to the boolean variables.) The replica formalism provides an order parameter describing the statistics of optimal assignments, that is, those ‘ground-state’ assignments that minimize the number of violated (unsatisfied) clauses. Consider an instance of the  $(2 + p)$ -SAT problem. We use the  $N_{GS}$  ground-state configurations to define  $m_i = 1/N_{GS} \sum_{g=1}^{N_{GS}} S_i^g$ , that is, the average value of spin  $S_i$  over all optimal configurations. Clearly,  $m_i$  ranges from  $-1$  to  $+1$ , and  $m_i = -1$  means that the corresponding boolean variable  $x_i$  is always false in all ground states; conversely,  $m_i = +1$  means that  $x_i$  is always true in all ground states. The distribution  $P(m)$  of all  $m_i$  characterizes the microscopic structure of the ground states. The accumulation of magnetizations  $m$  around  $\pm 1$  represents a ‘backbone’ of almost completely constrained variables, whose logical values cannot vary from solution to solution, while the centre of the distribution  $P(m \approx 0)$  describes weakly constrained variables.

### The backbone fraction is the order parameter

For  $\alpha < \alpha_c$ , the solution exhibits a simple symmetry property, usually referred to as replica symmetry, which leads to an order parameter which is precisely the magnetization distribution  $P(m)$  defined above. An essential qualitative difference between 2-SAT and 3-SAT is the way that the order parameter  $P(m)$  changes at the



**Figure 2** Theoretical and experimental results for the SAT/UNSAT transitions in the  $(2 + p)$ -SAT model. The vertical line at  $p_0$  ( $\sim 0.41$ ) separates the continuous from the discontinuous transition. The full line is the replica-symmetric theory’s predicted transition, believed exact for  $p < p_0$ . The diamond data points are results of computer experiment and finite-size scaling. The other two lines show upper and lower bounds obtained in ref. 21, while the stronger upper bound due to ref. 32 (see also refs 33, 34), and the best known lower bound, due to ref. 35, are indicated by the two square data points (only known for  $p = 1.0$ , that is,  $K = 3$ ). Inset, the values of  $\nu$  (the finite-size scaling exponent) obtained from finite-size scaling analysis. ( $K$  is the average number of variables per clause). We note that the finite-size scaling exponent  $\nu$  now depends on  $p$ , dropping roughly linearly to 1.5 at  $p = 1.0$ .

threshold. This discrepancy can be seen in the fraction  $f(K, \alpha)$  of boolean variables which become fully constrained, at and above the threshold. Below the threshold,  $f(K, \alpha) = 0$ . (Consider adding one clause to a satisfiable formula found below  $\alpha_c$ . If there is a finite fraction of backbone spins, there will be a finite probability that the added clause creates an unsatisfiable formula, but we know that the probability of an unsatisfiable formula tends to zero for large  $N$ , for  $\alpha < \alpha_c$ .) For 2-SAT,  $f(2, \alpha)$  becomes strictly positive above  $\alpha_c (= 1)$  and is continuous at the transition:  $f(2, 1^-) = f(2, 1^+) = 0$ . By contrast,  $f(3, \alpha)$  displays a discontinuous change at the threshold:  $f(3, \alpha_c^-) = 0$  and  $f(3, \alpha_c^+) = f_c(3) > 0$ .


For the mixed  $(2 + p)$ -SAT model, the crucial issue is therefore to understand how a discontinuous 3-SAT-like transition mechanism may appear when increasing  $p$  from zero up to one, and to see how it affects the computational cost of finding solutions near the threshold. Using replica methods, we find a continuous (second-order) SAT/UNSAT transition at  $\alpha_c(2 + p) = 1/(1 - p)$  for  $p < p_0$ , where  $p_0$ , estimated by various methods, lies between 0.4 and 0.416. This transition point can be calculated by simply ignoring the 3-SAT clauses. (Note that we have  $(1 - p)M$  2-SAT clauses, and for 2-SAT  $(M/N)_c = \alpha_c = 1$ .) It is perhaps somewhat counterintuitive that the 3-SAT clauses do not contribute at all to the location of the phase transition for  $p < p_0$  (in a sense, the 2-SAT constraints completely dominate the formulae), but this is consistent with a recent analytical result<sup>21</sup> derived for our model up to  $p = 0.4$ . The replica symmetric theory appears to be correct for all  $\alpha < \alpha_c(2 + p)$  when  $p < p_0$ , and predicts the threshold correctly, as in the  $K = 2$  case.

We have employed finite-size scaling analysis to locate the transition region<sup>8</sup>, in which the likelihood of a formula being UNSAT grows from zero to one, in a region centred around  $\alpha_c(2 + p)$  and of width proportional to  $N^{-1/\nu}$ . The values of  $\alpha_c$  obtained (Fig. 2) agree with the replica symmetric theory for  $p < p_0$ . Below  $p_0$ , the exponent  $\nu$  is roughly constant and close to 2.8, the value found for 2-SAT (Fig. 2, inset). Corrections to scaling could make  $\nu$  compatible with the value 3 obtained both from mean-field theory and from the recent exact bounds on the growth of the 2-SAT backbone (B. Bollobas, C. Borgs, J. Chayes, J. Kim, D. Wilson, personal communication). This suggests that the  $(2 + p)$ -SAT transition is in the same universality class as 2-SAT for  $p < p_0$ . So, for  $p < p_0$ , the  $(2 + p)$ -SAT model shares the physical features of random 2-SAT with a continuous phase transition. A plausible hypothesis is that  $\nu = 3$  for all  $p < p_0$ .

For  $p > p_0$ , the transition becomes discontinuous and the replica symmetric transition can only provide an upper bound for the true  $\alpha_c(2 + p)$ . The replica symmetric (RS) theory correctly predicts a discontinuous appearance of a finite fraction of fully constrained variables which jumps from 0 to  $f_c$  when crossing the threshold  $\alpha_c(2 + p)$ . However, values of both  $f_c(2 + p)$  and  $\alpha_c$  are overestimated; for example, for  $p = 1$ ,  $\alpha_c^{RS}(3) \approx 4.60$  and  $f_c^{RS}(3) \approx 0.94$ . Subtracting the terms in  $P(m)$  which appear to be splitting away from  $m = \pm 1$  reduces that to an estimated  $f_c(3) \approx 0.6$  in an approximation beyond replica symmetry, while experiments give  $\alpha_c(3) \approx 4.27$  and  $f_c(3) \approx 0.4$ . A replica symmetry breaking theory will be necessary to predict these quantities. For  $p > p_0$ , the random  $(2 + p)$ -SAT problem shares the physical features of random 3-SAT.

### Scaling of computational cost

Previous work showed that the cost of running the best heuristics, depth-first search algorithms with backtracking<sup>22</sup> increases exponentially with  $N$  for any value of  $\alpha \geq \alpha_c$  for  $K = 3$  (refs 7,23). The cost reaches its maximum value at  $\alpha_{0.5}(K, N)$ , the value of  $\alpha$  for which the probability of the formula being satisfiable equals 0.5, and we have the largest uncertainty about whether the formula is satisfiable or not, leading to the greatest difficulty for the search algorithm. Our data show that the computational cost at  $\alpha_{0.5}(K, N)$



**Figure 3** The median computation cost (number of backtracks) of proving a formula SAT or UNSAT in the  $(2 + p)$ -SAT model for a range of values of  $p$ . We used a state-of-the-art implementation of the Davis–Putnam search method<sup>36</sup> for  $p = 0.0, 0.2, 0.4, 0.6$ , fixing  $\alpha = \alpha_{0.5}$  (the hardest region), over the range of  $N$  that could be efficiently searched. The semi-log plot shows an exponential dependence of cost on  $N$  for  $p = 0.6$  ( $p > p_0$ ), but it remains linear in  $N$  for  $p \leq p_0$ .


scales linearly with  $N$  for  $p = 0$  through  $0.4$ , and exponentially (Fig. 3) for  $p = 0.6$ . These results are again consistent with the observation that  $(2 + p)$ -SAT behaves essentially as 2-SAT for  $p \leq p_0$ , and as 3-SAT for  $p > p_0$ .

The presence of a ‘backbone’ of frozen spins provides a good explanation for the apparently high cost of heuristic search near the phase boundary. A backtrack search method which proceeds by asserting possible spin values can make early mistakes by setting backbone spins (variables) to incorrect values, and then take a long time to backtrack to correct these mistakes. (The search may have to explore an exponential size subtree with no solutions<sup>24</sup>.) The positive value of  $f(2, \alpha)$  above  $\alpha_c$  is consistent with the fact that finding the ‘best’ assignment, that is, one that satisfies the most clauses, is NP-complete for 2-SAT.

Our experiments confirm that the appearance of the backbone is discontinuous (Fig. 4). Above the threshold, the fraction of frozen spins which was found in small samples by exhaustive enumeration of all ground states is relatively insensitive to  $N$ . Below the threshold, the fraction of frozen spins decreases rapidly with increasing  $N$ . Although the samples that could be studied are small, the results suggest that  $f(2 + p, \alpha)$  vanishes below  $\alpha_c$  and are consistent with  $f(2, \alpha_c^+)$  vanishing while  $f(3, \alpha_c^+)$  reaches a limit of slightly more than  $0.4$  for large  $N$ .

**Discussion**

The random first-order phase transition in  $K$ -SAT is unusual because it has a mixed nature: of first order in the order parameter, yet continuous in thermodynamic quantities such as the entropy<sup>12,25</sup>. This means that while an extensive number of degrees of freedom freeze at the transition, forming the backbone, the remaining ones can support critical fluctuations. The separation of the backbone is a realization of the ‘two components’ (corresponding to slow and fast degrees of freedom) idea sometimes put forward on heuristic grounds to describe the physics of glasses<sup>26</sup> or granular materials<sup>27</sup>. The exponent  $\nu$  shown in Fig. 2 inset is quite different from the value,  $1$ , expected for typical first-order transitions, or for the mean-field theories that usually describe models with no spatial dimensionality.



**Figure 4** Backbone fractions as a function of  $\alpha$  for 2-SAT and 3-SAT. The data were extracted from 3,000–15,000 samples for the 2-SAT cases (on the left) with  $N$  values of 20 (red), 30 (blue), 45 (green), 100 (black), 200 (light blue), and 500 (orange). The results for  $N \leq 45$  were obtained by exhaustive enumeration, examining all assignments, while those for  $N \geq 100$  used a modified Davis–Putnam search procedure. The 3-SAT cases were studied by exhaustive enumeration in 7,500–30,000 samples for  $N$  values of 16 (red), 20 (blue), 24 (green), and 28 (black). The vertical lines mark the observed SAT/UNSAT thresholds in the limit  $N \rightarrow \infty$ . For 2-SAT, data obtained from larger sizes show that the backbone fraction at the threshold tends towards zero.

The clear experimental evidence for the coexistence of a first-order transition with critical fluctuations is, to our knowledge, new with  $K$ -SAT, although there have been hints of such a possibility occurring in ‘rigidity percolation’ models<sup>28,29</sup> intended to capture the mechanical properties of sand and similar granular materials.

We note that conventional first-order transitions, such as the freezing of a solid from its liquid form, do not imply computational complexity. Describing the ground state of a crystal, once the positions of a few atoms are established, requires only the application of symmetry to calculate positions for the remaining atoms with an effort that would be linear in the number of atoms. By contrast, a spin glass may have an infinite number of ground states (in the limit  $N \rightarrow \infty$ ), and no translational symmetry.

The existence of a backbone has also been observed in other combinatorial problems, such as the travelling salesman problem<sup>30</sup>. It is possible to improve heuristic optimization methods, at least for this problem, by early identification and elimination of the backbone. The remaining problem consists of many separate subproblems, each of which is much less complex to solve<sup>31</sup>. This may prove to be a generally valid approach. Efficient means of finding the backbone would be specific to each problem type, but should nonetheless provide a step forwards in algorithm efficiency. Moreover, many NP-complete problems occurring in practice contain a mix of tractable and intractable constraints. Our results suggest that the run time of search algorithms that exploit as much of the tractable structure as possible may in fact scale polynomially in the typical case. We hope that our work will stimulate other studies to develop a further understanding of the structure of phase transitions taking place in computational problems, as studied in computer science, and in diluted random systems, as studied in statistical physics. Both fields seem to share several basic open problems. □

Received 21 December 1998; accepted 11 May 1999.

1. Cook, S. A. in *Proc. 3rd Annu. ACM Symp. on Theory of Computing* 151–158 (Assoc. Comput. Machinery, New York, 1971).
2. Garey, M. & Johnson, D. S. *Computers and Intractability: A Guide to the Theory of NP-completeness* (Freeman, San Francisco, 1979).
3. Papadimitriou, C. *Computational Complexity* (Addison-Wesley, Reading, 1994).

4. Fu, Y. T. & Anderson, P. W. Application of statistical mechanics to NP-complete problems in combinatorial optimization. *J. Phys. A* **19**, 1605–1620 (1985).
5. Mézard, M., Parisi, G. & Virasoro, M. A. *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
6. Cheeseman, P., Kanefsky, B. & Taylor, W. M. in *Proc. 13th Int. Joint Conf. on Artificial Intelligence (IJCAI-91)* (eds Mylopoulos, J. & Reiter, R.) 331–337 (Morgan Kaufmann, San Mateo, California, 1991).
7. Mitchell, D., Selman, B. & Levesque, H. in *Proc. 10th Natl. Conf. on Artificial Intelligence (AAAI-92)* 440–446 (AAAI Press/MIT Press, Cambridge, Massachusetts, 1992).
8. Kirkpatrick, S. & Selman, B. Critical behavior in the satisfiability of random Boolean expressions. *Science* **264**, 1297–1301 (1994).
9. Selman, B., Mitchell, D. & Levesque, H. Generating hard satisfiability problems. *Artif. Intell.* **81**, 17–29 (1996).
10. Hogg, T., Huberman, B. A. & Williams, C. (eds) *Frontiers in problem solving: phase transitions and complexity*. *Artif. Intell.* **81**(I & II) (1996).
11. Hayes, B. Can't get no satisfaction. *Am. Sci.* **85**, 108–112 (1996).
12. Gross, D. & Mezard, M. The simplest spin glass. *Nucl. Phys. B* **240**, 431–452 (1984).
13. Aspvall, B., Plass, M. F. & Tarjan, R. E. A linear-time algorithm for testing the truth of certain quantified Boolean formulas. *Inf. Process. Lett.* **8**, 121–123 (1979).
14. Goerdts, A. A threshold for unsatisfiability. *J. Comput. System Sci.* **53**, 469–486 (1996).
15. Chvátal, V. & Reed, B. Mick gets some (the odds are on his side). *Proc. 33rd IEEE Symp. on Foundations of Computer Science* 620–627 (IEEE Computer Soc. Press, New York, 1992).
16. Friedgut, E. Sharp thresholds of graph properties and the K-SAT problem. *J. Am. Math. Soc.* (in the press).
17. Beame, P., Karp, R., Pitassi, T. & Saks, M. in *Proc. ACM Symp. on Theory of Computing (STOC98)* 561–571 (Assoc. Comput. Machinery, New York, 1998).
18. Monasson, R. & Zecchina, R. Entropy of the K-satisfiability problem. *Phys. Rev. Lett.* **76**, 3881–3885 (1996).
19. Monasson, R. & Zecchina, R. The statistical mechanics of the random K-satisfiability model. *Phys. Rev. E* **56**, 1357–1370 (1997).
20. Monasson, R. & Zecchina, R. Tricritical points in random combinatorics: the 2+p SAT case. *J. Phys. A* **31**, 9209–9217 (1998).
21. Achlioptas, D., Kirousis, L., Kranakis, E. & Krizanc, D. in *Proc. Workshop on Randomized Algorithms in Sequential, Parallel and Distributed Computing (RALCOM 97)*, 1–10 (1998).
22. Davis, M. & Putnam, H. A computing procedure for quantification theory. *J. Assoc. Comput. Machinery* **7**, 201–215 (1960).
23. Selman, B. & Kirkpatrick, S. Critical behavior in the cost of K-satisfiability. *Artif. Intell.* **81**, 273–295 (1996).
24. Gomes, C. P., Selman, B. & Crato, N. in *Principles and Practice of Constraint Programming (CP97)* (ed. Smolka, G.) 121–135 (Lect. Notes in Comput. Sci., Vol. 1330, Springer, Heidelberg, 1997).
25. Kirkpatrick, T. R., Thirumalai, D. & Wolynes, P. G. Scaling concepts for the dynamics of viscous liquids near an ideal glassy state. *Phys. Rev. A* **40**, 1045–1054 (1989).
26. Angell, A. C. Formation of glasses from liquids and biopolymers. *Science* **267**, 1924–1935 (1995).
27. Bouchaud, J.-P., Cates, M. E., Ravi Prakash, J. & Edwards, S. F. Hysteresis and metastability in a continuum sandpile model. *Phys. Rev. Lett.* **74**, 1982–1985 (1995).
28. Jacobs, D. J. & Thorpe, M. F. Generic rigidity percolation: the pebble game. *Phys. Rev. Lett.* **75**, 4051–4054 (1995).
29. Moukarzel, C. & Duxbury, P. M. Stressed backbone and elasticity of random central-force systems. *Phys. Rev. Lett.* **75**, 4055–4058 (1995).
30. Kirkpatrick, S. & Toulouse, G. Configuration space analysis of traveling salesman problems. *J. Phys. (France)* **46**, 1277–1292 (1985).
31. Schneider, J., Froschhammer, C., Morgenstern, I., Husslein, T. & Singer, J. M. Searching for backbones—an efficient parallel algorithm for the traveling salesman problem. *Comput. Phys. Commun.* **96**, 173–188 (1996).
32. Kirousis, L., Kranakis, E. & Krizanc, D. in *Proc. 4th European Symp. on Algorithms* 27–38 (Springer, Berlin, 1996).
33. Dubois, O. & Boufkhad, Y. A general upper bound for the satisfiability threshold of random K-SAT formulas. *J. Algorithms* **24**, 395–420 (1997).
34. Kamath, A., Motwani, R., Palem, K. & Spirakis, P. Tail bounds for occupancy and the satisfiability threshold conjecture. *Random Struct. Algorithms* **7**, 59–89 (1995).
35. Frieze, A. & Suen, S. Analysis of two simple heuristics on a random instance of K-SAT. *J. Algorithms* **20**, 312–335 (1996).
36. Crawford, J. & Auton, L. in *Proc. 11th Natl. Conference on Artificial Intelligence (AAAI-93)* 21–27 (AAAI Press/MIT Press, Cambridge, Massachusetts, 1993).

**Acknowledgements.** S.K. thanks the Laboratories of Theoretical and Statistical Physics at Ecole Normale Supérieure for a visiting professorship in 1998, during which this work was completed. B.S. is an Alfred P. Sloan research fellow and is supported by an NSF Faculty Early Career Development Award. During the early phase of this research, B.S. was at AT&T Bell Laboratories.

Correspondence and requests for materials should be addressed to S.K. (e-mail: kirk@watson.ibm.com).