Circumspect descent prevails in solving random constraint satisfaction problems

Mikko Alava*, John Ardelius†‡, Erik Aurell‡§¶, Petteri Kaski‖, Supriya Krishnamurthy†, Pekka Orponen††, and Sakari Seitz*††**

*Department of Engineering Physics, Helsinki University of Technology, P.O. Box 1100, FI-02015, Espoo, Finland; †Swedish Institute of Computer Science AB, SE-164 29 Kista, Sweden; ‡Department of Computational Biology, AlbaNova University Centre, SE-106 91 Stockholm, Sweden; §Linnaeus Centre, KTH-Royal Institute of Technology, SE-100 44 Stockholm, Sweden; ‖Helsinki Institute for Information Technology (HIIT), Department of Computer Science, University of Helsinki, P.O.Box 68, FI-00014, Helsinki, Finland; **School of Information and Communication Technology, KTH-Royal Institute of Technology, SE-164 40 Kista, Sweden; and ††Department of Information and Computer Science, Helsinki University of Technology, P.O. Box 5400, FI-02015, Espoo, Finland

Edited by Giorgio Parisi, University of Rome, Rome, Italy, and approved August 22, 2008 (received for review January 16, 2008)

We study the performance of stochastic local search algorithms for random instances of the *K*-satisfiability (*K*-SAT) problem. We present a stochastic local search algorithm, ChainSAT, which moves in the energy landscape of a problem instance by never going upwards in energy. ChainSAT is a focused algorithm in the sense that it focuses on variables occurring in unsatisfied clauses. We show by extensive numerical investigations that ChainSAT and other focused algorithms solve large *K*-SAT instances almost surely in linear time, up to high clause-to-variable ratios α; for example, for $K = 4$ we observe linear-time performance well beyond the recently postulated clustering and condensation transitions in the solution space. The performance of ChainSAT is a surprise given that by design the algorithm gets trapped into the first local energy minimum it encounters, yet no such minima are encountered. We also study the geometry of the solution space as accessed by stochastic local search algorithms.

geometry of solutions | local search | performance | random K-SAT

COMPUTER SCIENCES

Constraint satisfaction problems (CSPs) are the industrial, commercial, and often very large-scale analogues of popular leisure-time pursuits such as the Sudoku puzzle. They can be formulated abstractly in terms of N variables $x_1, x_2, \ldots, x_N$ and M constraints, where each variable x_i takes a value in a finite set, and each constraint forbids certain combinations of values to the variables. The classical example of a worst-case intractable (1) constraint satisfaction problem is the K-satisfiability (K-SAT) problem (2), where each variable takes a Boolean value (either 0 or 1), and each constraint is a clause over K variables disallowing one out of the 2^K possible combinations of values. An instance of K-SAT can also be interpreted directly as a spin system of statistical physics. Each constraint equals to a K-spin interaction in a Hamiltonian, and thus spins represent the original variables; ground states of the Hamiltonian at zero energy correspond to the solutions, that is, assignments of values to the variables that satisfy all of the clauses (3).

It was first observed in the context of K-SAT, and then in the context of several other CSPs (4), that ensembles of random CSPs have a "phase transition," a sharp change in the likelihood to be solvable (5). Empirically, algorithms have been observed to fail or have difficulties in the immediate neighborhood of such phase-transition points, a fact that has given rise to a large literature (4). Large unstructured CSPs are solved either by general-purpose deterministic methods, of which the archetypal example is the Davis–Putnam–Logemann–Loveland (DPLL) algorithm (6) or using more tailored algorithms, such as the Survey Propagation (SP) algorithm (7) motivated by spin glass theory, or variants of stochastic local search techniques (8–10).

Stochastic local search (SLS) methods are competitive on some of the largest and least-structured problems of interest (11), in particular on random K-SAT instances, which are constructed by selecting independently and uniformly at random M clauses over the N variables, where the parameter controlling the satisfiability of an instance is $\alpha = M/N$, the ratio of clauses to variables. SLS algorithms work by making successive random changes to a trial configuration (assignment of values to the variables) based on information about a local neighborhood in the set of all possible configurations. Their modern history starts with the celebrated simulated annealing algorithm of Kirkpatrick, Gelatt, and Vecchi (12). From the perspective of K-SAT, the next fundamental step forward was an algorithm of Papadimitriou (13), now often called RandomWalkSAT, which introduced the notion of focusing the random moves to rectify broken constraints. RandomWalkSAT has been shown, by simulation and theoretical arguments, to solve the paradigmatic case of random 3-satisfiability up to about $\alpha = 2.7$ clauses per variable, almost surely in time linear in N (14, 15). A subsequent influential development occurred with Selman, Kautz, and Cohen's WalkSAT algorithm (16), which mixes focused random and greedy moves for better performance. We have previously shown that WalkSAT and several other stochastic local search heuristics work almost surely in linear time, up to at least $\alpha = 4.21$ clauses per variable (17–19). In comparison, the satisfiability/unsatisfiability threshold of random 3-satisfiability is believed to be at $\alpha = 4.267$ clauses per variable (20).

The present work carries out a systematic empirical study of random K-SAT for $K = 4$, and we also present extensive data for $K = 5$ and $K = 6$. Our motivation for this study is 3-fold.

Testing the Limits of Local Search. It has been empirically observed for $K = 3$ that many SLS algorithms have a linear-time regime, which extends to the immediate vicinity of the phase transition point (17–19). Thus, a similar investigation for higher K is warranted.

Structure of the Space of Solutions. Recent rigorous results and nonrigorous predictions from spin-glass theory suggest that the structure of the space of solutions of a random K-SAT instance undergoes various qualitative changes for $K \geq 4$, the implications of which to the performance of algorithms should be investigated.

Mézard, Mora, and Zecchina (21) have shown rigorously that for $K \geq 8$, the space of solutions of random K-SAT breaks into multiple clusters separated by extensive Hamming distance. (The Hamming distance of two Boolean vectors of length N is the number positions in which the vectors differ divided by N.) In more precise terms, an instance of K-SAT is *x-satisfiable* if it has

Author contributions: M.A. and E.A. designed research; M.A., J.A., E.A., P.K., and S.S. performed research; S.S. contributed new reagents/analytic tools; M.A., E.A., P.K., S.K., and S.S. analyzed data; and M.A., E.A., P.K., and P.O. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¶To whom correspondence should be addressed. E-mail: eaurell@kth.se.

a pair of solutions with normalized Hamming distance $0 \le x \le 1$. Mézard, Mora, and Zecchina (21) show that, for $K \ge 8$, there exists an interval (a, b), $0 < a < b < 1/2$, such that, with high probability as $N \to \infty$, a random instance ceases to be x-satisfiable for all $x \in (a, b)$ at a smaller value of α before it ceases to be x-satisfiable for some $x \in [b, 1/2]$.

For $K = 4$, we see no evidence of gaps in the empirical x-satisfiability spectrum in the linear-time regime of SLS algorithms, which includes the predicted spin-glass theoretic clustering points. In light of the rigorous results for $K \ge 8$, this suggests that the cases $K = 4$ and $K = 8$ may be qualitatively different. Moreover, we observe that recently predicted spin-glass-theoretic clustering thresholds [Krzakala *et al.* (22)] have no impact on algorithm performance. This puts forth the question whether the energy landscape of random K-SAT for small K is in some regard more elementary than has been previously believed.

Structure of the Energy Landscape. In the context of random K-SAT, it is folklore that SLS algorithms appear to benefit from circumspect descent in energy, that is, from a conservative policy of lowering the number of clauses not satisfied by the trial configuration. To explore this issue further, we introduce a SLS algorithm which we call ChainSAT. It is based on four ideas: (*i*) focusing, (*ii*) easing difficult-to-satisfy constraints by so-called chaining moves, (*iii*) restraining the rate of descent, and (*iv*) never going upward in energy; that is, the number of unsatisfied clauses is a nonincreasing function of the sequence of trial configurations traversed by the algorithm.

By design, ChainSAT cannot escape from a local minimum of energy in the energy landscape. Yet, empirically, ChainSAT is able to find (for the $K = 4, 5, 6$ studied here) a solution, almost surely in linear time, up to values of α reached by SLS algorithms that are allowed to go up in energy, such as the Focused Metropolis Search (19). This observation further supports the position that random K-SAT for small K may be more elementary than has been previously believed.

Results

Experiments with Focused Metropolis Search. The Focused Metropolis Search (FMS) algorithm (19) is given in pseudocode as follows:

```
S = random assignment of values to the variables
while S is not a solution do
    C = a clause not satisfied by S selected uniformly at
        random
    V = a variable in C selected uniformly at random
    ΔE = change in number of unsatisfied clauses if V is
         flipped in S
    if ΔE ≤ 0 then
       flip V in S
    else
       with probability η^ΔE
          flip V in S
       end with
    end if
end while
```

This section documents our experiments aimed at charting the empirical linear-time region of FMS on random K-SAT for $K = 4$.

FMS Performance. In Fig. 1, we present empirical evidence that FMS almost surely runs in time linear in N for instances of random K-satisfiability at $K = 4$ and $\alpha = 9.6$. That the curves get steeper with increasing N implies concentration of solution times, or that above- and below-average solution times get rarer with increasing N. Note that the scaling implies performance almost surely linear in N and demonstrates that the linear-time regime of FMS extends beyond the predicted (22) spin-glass theoretic "dynamical" and "condensation" transitions points.

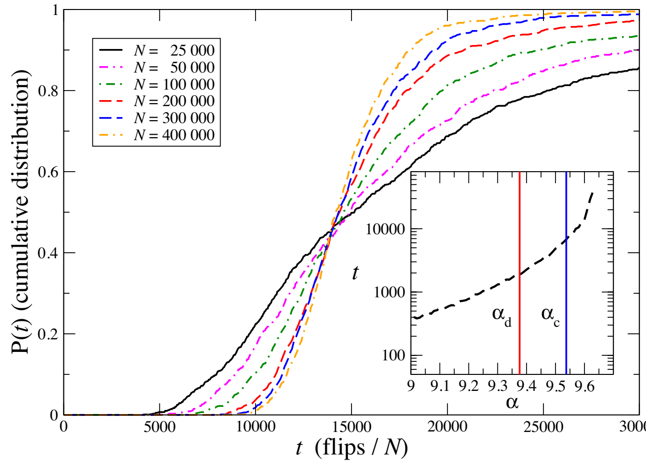

Fig. 1. Cumulative distributions of solution times normalized by the number of variables N for the Focused Metropolis Search algorithm (19) on instances of random K-satisfiability at $K = 4$ and $\alpha = 9.6$. The vertical axis indicates the fraction of 1001 random instances solved within a given running time, measured in flips / N on the horizontal axis. (*Inset*) Here, we present the scaling of the algorithm as α increases (with $N = 100000$). The "temperature" parameter of FMS is set to $\eta = 0.293$.

For $K = 3$, it has already been established that the FMS algorithm has an "operating window" in terms of the adjustable "temperature" parameter η (19). For too-large values of η, the linearity (in N) is destroyed due to too-large fluctuations that keep the algorithm from reaching low energies and the solution. For too small values of η, the algorithm becomes "too greedy," leading to a divergence of solution times. Thus, to obtain performance linear in N, it is necessary to carefully optimize the parameter η. Such an optimization for $K = 4$ (the details of which we omit for reasons of space; compare ref. 19) reveals that the operating window for $\alpha = 9.60$ is at least $0.292 \le \eta \le 0.294$. In the experiments, we have chosen $\eta = 0.293$.

Experiments on x-Satisfiability Using FMS. Our experimental setup to investigate x-satisfiability is as follows. For given values of α and N, we first generate a random K-SAT instance and find one reference solution of this instance using FMS. Then, using FMS, we search for other solutions in the same instance. The initial configuration S for FMS is selected uniformly at random from the set of all configurations having a given Hamming distance to the reference solution. When FMS finds a solution, we record the distance x of the solution found to the reference solution.

Our experiments on random K-SAT for $K = 4$ did not reveal any gaps in the x-satisfiability spectrum, even for $\alpha = 9.6$, beyond the predicted spin-glass theoretic "dynamical" and "condensation" transitions points (22). In particular, Fig. 2 gives empirical evidence that solutions are found at all distances smaller than the typical distance of solutions found by FMS. This is in contrast to the numerical results of Battaglia *et al.* (23) for a balanced version of $K = 5$.

Here, it should be pointed out that the solutions found by stochastic local search need not be typical solutions in the space of all solutions; there can be other solutions that are not reached by FMS or other algorithms. Evidence of this is reflected in the "whiteness" status of solutions (19, 24); all of the solutions found in our experiments were completely white, that is, they do not have locally frozen variables (25). One can, of course, imagine that a "typical solution" is not white under the circumstances

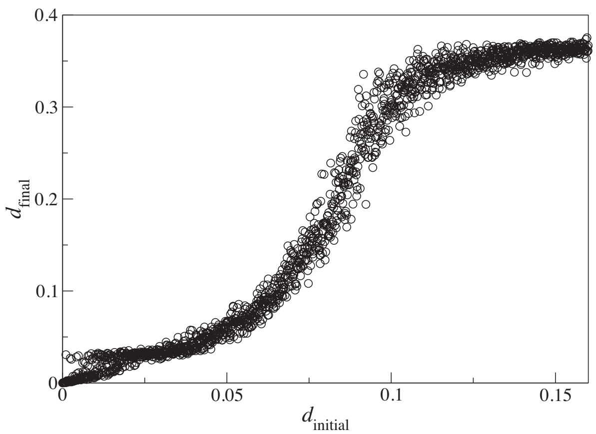

Fig. 2. Investigation of x-satisfiability using FMS initialized with a random configuration at a given Hamming distance from a reference solution. One reference solution and one instance of random K-SAT at $K = 4$, $\alpha = 9.6$, and $N = 200000$. The horizontal axis displays the normalized Hamming distance of the initial configuration to the reference solution. The vertical axis displays the normalized Hamming distance of the solution found to the reference solution. All plotted 1601 searches produced a solution, and no gaps are visible in the vertical axis, suggesting asymptotic x-satisfiability for $x \leq 0.37$. The temperature parameter of FMS is set to $\eta = 0.293$.

examined here, but, as noted, there is no evidence of the existence of such (compare ref. 26). Fig. 3 summarizes the results of a scaling analysis with increasing N over five random instances and reference solutions. The distance distributions appear to converge to some specific curve without vertical sections, the absence of which suggests that the x-satisfiability spectrum has no gaps below the typical distance of solutions found by FMS in the limit of infinite N.

Fig. 4 summarizes the results of a scaling analysis with increasing α. We see that the typical distance between solutions found by FMS decreases with increasing α, and that no clear gaps are apparent in the distance data.

Experiments with ChainSAT. A heuristic that never moves up in energy is here shown to solve random K-satisfiability problems,

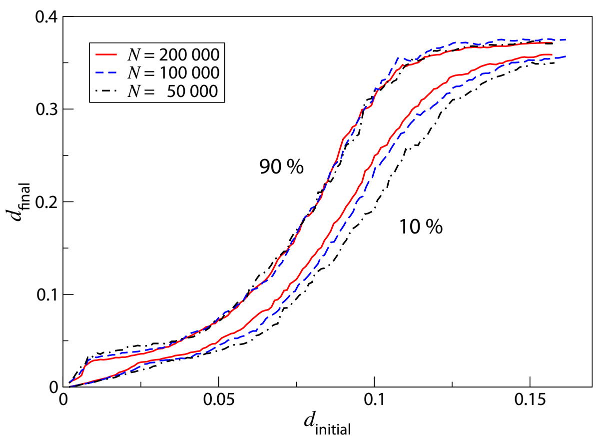

Fig. 3. Scaling of x-satisfiability data obtained using FMS on random K-SAT with increasing N. The parameters $K = 4$, $\alpha = 9.6$, and $\eta = 0.293$ are fixed. The plotted 10- and 90-percentile curves are calculated from five random instances and reference solutions for each $N = 50000$, 100000, 200000, with a moving window size of 0.004 in the horizontal axis. The distances appear to converge close to the 90-percentile curves.

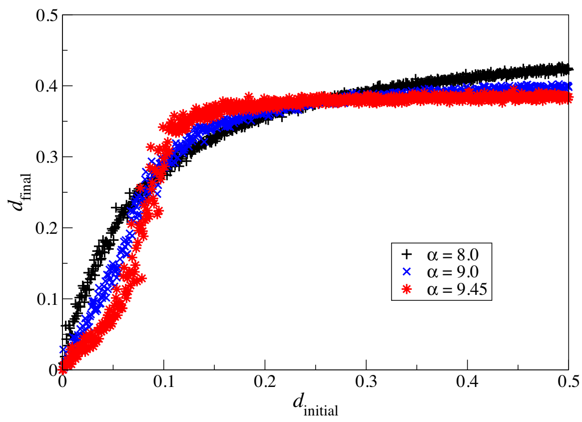

Fig. 4. Scaling of x-satisfiability data obtained using FMS on random K-SAT with increasing α. The parameters $K = 4$, $N = 100000$, and $\eta = 0.293$ are fixed. One random instance and one reference solution for each $\alpha = 8.0$, 9.0, 9.45; see Fig. 2 for $\alpha = 9.6$. The value $\alpha = 9.45$ is between the predicted locations of the dynamical and the condensation transition points (22). No clear gap in distances is discernible in any of the cases.

almost surely in time linear in N, for $K = 4, 5, 6$. Our heuristic, ChainSAT, is given in pseudocode as follows:

S = random assignment of values to the variables
chaining = FALSE
while S is not a solution **do**
  if not chaining **then**
    C = a clause not satisfied by S selected uniformly at random
    V = a variable in C selected u.a.r.
  end if
  ΔE = change in number of unsatisfied clauses if V is flipped in S
  chaining = FALSE
  if $\Delta E = 0$ **then**
    flip V in S
  else if $\Delta E < 0$
    with probability p_1
      flip V in S
    end with
  else
    with probability $1 - p_2$
      C = a clause satisfied only by V selected u.a.r.
      V' = a variable in C other than V selected u.a.r.
      $V = V'$
      chaining = TRUE
    end with
  end if
end while

The algorithm (i) never increases the energy of the current configuration S and (ii) exercises circumspection in decreasing the energy. In particular, moves that decrease the energy are taken only sporadically compared with equienergetic and chaining moves. The latter are designed to alleviate critically satisfied constraints by proceeding in "chains" of variable-clause-variable until a variable is found that can be flipped without increase in energy. Focusing is used for the nonchaining moves. The structure of ChainSAT has the basic idea of helping to flip a variable to satisfy an original broken constraint.

The ChainSAT algorithm has two adjustable parameters, one (p_1) for controlling the rate of descent (by accepting energy-lowering flips) and another (p_2) for limiting the length of the chains to avoid looping. In our experiments, we let $p = p_1 = p_2$.

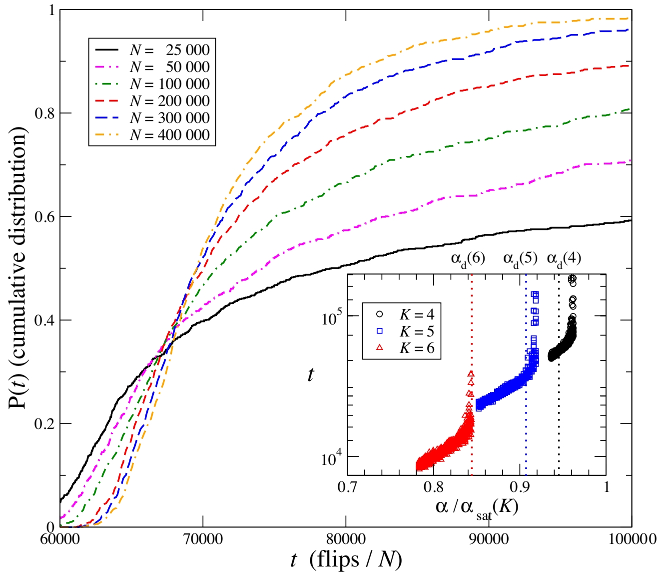

Fig. 5. Cumulative distributions of solution times normalized by number of variables N for the ChainSAT algorithm on random K-satisfiability instances at $K = 4$ and $\alpha = 9.55$. The vertical axis indicates the fraction of 1001 random input instances solved within a given running time, measured in flips / N on the horizontal axis. (*Inset*) Here, we present the scaling of the algorithm for $K = 4, 5, 6$ at $N = 100000$ with increasing α; the values of $\alpha(K)$ in the horizontal axis have been normalized with $\alpha_{sat}(K)$, which has the population dynamics estimated values $\alpha_{sat}(4) = 9.931$, $\alpha_{sat}(5) = 21.117$, and $\alpha_{sat}(6) = 43.37$ (20). The parameters of ChainSAT have been chosen to be small enough to work at least up to the predicted "dynamical transition" (22): we have set $p_1 = p_2 = 0.0001$ ($K = 4$), 0.0002 ($K = 5$), and 0.0005 ($K = 6$).

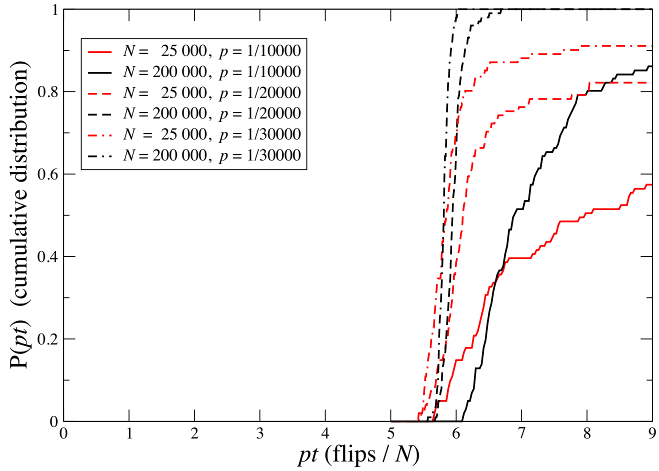

Fig. 6. Sensitivity of the ChainSAT algorithm to the parameter $p = p_1 = p_2$ on random K-satisfiability instances at $K = 4$ and $\alpha = 9.55$. Cumulative distributions of solution times for $1 / p = 10000, 20000, 30000$ at $N = 25000, 200000$. The vertical axis indicates the fraction of 101 random instances solved within a given running time, measured in $p \times$ flips / N on the horizontal axis.

ChainSAT Performance. In Fig. 5, we present empirical evidence that ChainSAT almost surely runs in time linear in N for random K-satisfiability problems with $K = 4, 5, 6$. That the curves get steeper with increasing N implies concentration of solution times, or that above-and below-average solution times get rarer with N. Because the algorithm never goes uphill in the energy landscape, local energy minima cannot be an obstruction to finding solutions, at least in the region of the energy landscape visited by this algorithm. On the other hand, when ChainSAT fails to find a solution in linear time, this can also result from simply getting lost; in particular, the fraction of moves that lower the energy over those that keep it constant may dwindle to zero.

Fig. 6 illustrates the sensitivity of ChainSAT to the choice of the parameter p as N is increased on random K-satisfiability problems with $K = 4$. We observe that as p decreases, the solution times concentrate more rapidly about the mean with increasing N. Furthermore, the mean solution time scales roughly as $1 / p$ when p is decreased.

Whiteness. To provide a further empirical analysis of ChainSAT, we next present Fig. 7. This is discussed not in terms of solution times and the range of α achievable with a bit of tuning, but in terms of two quantities: (*i*) the average chain length l_{chain} during the course of finding a solution and (*ii*) the average whiteness depth (AWD). In more precise terms, the average chain length is $l_{\mathrm{chain}} = f / m - 1$, where f is the total number of iterations of the main loop of ChainSAT and m is the number of times the if-statement controlled by the chaining flag in the main loop is executed.

The AWD is related to the result of the so-called whitening procedure (24), described in pseudocode as follows:

```
initially all clauses and variables are unmarked (non-white)
mark (whiten) every clause that is unsatisfied
mark (whiten) every clause that has more than one true literal
D = 0
repeat
    mark (whiten) any unmarked variables that appear as
    satisfying literals only in marked clauses
    if all the variables are marked then
        declare that S is completely white
        halt
    end if
    if no new variables were marked in this iteration then
        declare that S has a core
        halt
    end if
    mark (whiten) any unmarked clauses that contain at
    least one marked variable
    D = D + 1
end repeat
```

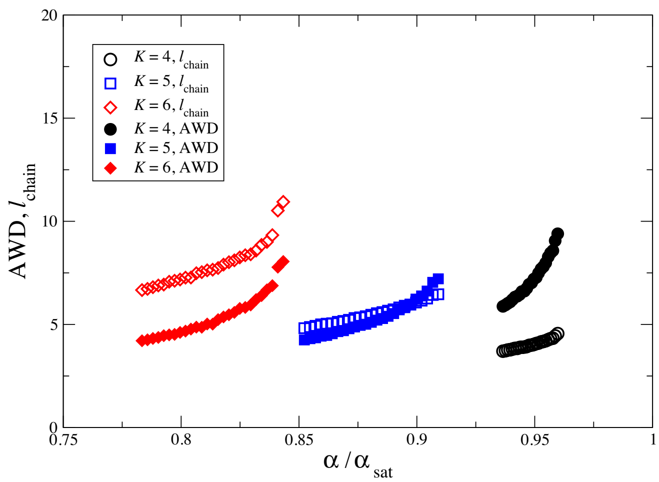

Fig. 7. The average chain length in ChainSAT and the average whiteness depth of the solutions found in random K-SAT for $K = 4, 5, 6$. Each plotted value is the average over 21 random instances. The values of $\alpha(K)$ in the horizontal axis have been normalized with $\alpha_{sat}(K)$, which has the empirical values $\alpha_{sat}(4) = 9.931$, $\alpha_{sat}(5) = 21.117$, and $\alpha_{sat}(6) = 43.37$ (20). The ChainSAT parameters are set to $p_1 = p_2 = 0.0001$ ($K = 4$), 0.0002 ($K = 5$), and 0.0005 ($K = 6$).

The whitening algorithm is applied to the solution found when ChainSAT terminates. The whiteness depth of a variable is defined as the value of D in the whitening procedure at the time the variable gets marked (whitened); the value is infinite if the variable never gets marked (whitened) during the whitening procedure. The AWD of a solution is the average of the whiteness depths of the variables. See ref. 19 for an empirical discussion of whitening in the context of random K-SAT for $K = 3$. The key observation here is that the solutions found by ChainSAT all have a finite AWD. This in loose terms means that there is "slack" in the solution.

Based on Fig. 7, it is clear that increasing the value of α has the same effect for $K = 4, 5, 6$: the average chain length l_{chain} increases and so does the AWD. Note that the ratio AWD / l_{chain} increases with α.

Discussion

We have here shown empirically that local search heuristics can be designed to avoid traps and "freezing" in random K-satisfiability, with solution times scaling linearly in N. This requires that circumspection is exercised; too greedy a descent causes the studied algorithms to fail for reasons unclear. A physics-inspired interpretation is that during a run the algorithm has to "equilibrate" on a constant energy surface.

In terms of the parameter α, it is the pertinent question as to how far the "easy" region from which one finds these solutions extends. For small K, it may be possible that this is true all the way to the satisfiability/unsatisfiability transition point. The empirical evidence we have here presented points toward a divergence of the prefactor of the linear scaling in problem size well below α_{sat}. Furthermore, this divergence is stronger for higher values of K. For large values of K, the absence of traps may, however, in any case be considered unlikely, as the rigorous techniques used to show clustering of solutions for $K \geq 8$ (21) can also be used to show that there exist pairs of distant solutions separated by an extensive energy barrier from each other. This suggests also the existence of local minima separated by extensive barriers. On the other hand, our present results for small K give no evidence in this direction. In particular, for $K = 4$, we have shown empirically that the energy landscapes can be navigated with simple randomized heuristics beyond all so far predicted transition points, apart from the satisfiability/unsatisfiability transition itself.

Our experiments also strongly suggest that the space of solutions for $K = 4$ at least up to $\alpha = 9.6$ does not break into multiple clusters separated by extensive distance. All of the solutions found have "slack," in the sense that they have a finite AWD. Is there an efficient way to find solutions that are not "white" in this sense; put otherwise, is the existence of "white" solutions necessary for "easy" solvability?

The observed success of ChainSAT adds evidence to the long-held belief in computer science that high-dimensional search spaces rarely have true local minima. This presents a contrast with the common practice to attribute the effectiveness of methods such as simulated annealing to the method's ability to make "uphill" moves. Our experiments suggest that "horizontal" moves ($\Delta E = 0$) are equally attributable.

All these observations present further questions about the structure of the energy landscape, the solution space, and the workings of algorithms for random CSPs. They also leave us with challenges and constraints to theoretical attempts to understand these, including approaches from the physics of spin glasses.

ACKNOWLEDGMENTS. This work was supported by the Integrated Project EVERGROW of the European Union (J.A. and S.K.), the Swedish Science Council through Linnaeus Centre ACCESS (E.A.), and the Academy of Finland (Grant 117499, to P.K.), and through the Center of Excellence Program (M.A. and S.S.). We thank the Department of Computer Science of University of Helsinki and Swedish Institute of Computer Science for the use of a little over 13 years of central processing unit time, and the Kavli Institute for Theoretical Physics China (KITPC) for hospitality.

COMPUTER SCIENCES

1. Garey MR, Johnson DS (1979) *Computers and Intractability: A Guide to the Theory of NP-Completeness* (Freeman, San Francisco).
2. Du D, Gu J, Pardalos P, eds (1997) *Satisfiability Problem: Theory and Applicationsm,* DIMACS Series in Discrete Mathematics and Theoretical Computer Science (Am Math Soc, Providence RI), Vol 35.
3. Martin OC, Monasson R, Zecchina R (2001) Statistical mechanics methods and phase transitions in optimization problems. *Theor Comput Sci* 265:3–67.
4. Dubois O, Monasson R, Selman B, Zecchina R, eds (2001) Special Issue on NP-Hardness and Phase Transitions. *Theor Comput Sci* 265.
5. Mitchell D, Selman B, Levesque H (1992) *Hard and Easy Distributions of SAT Problems*, Proceedings of the 10th National Conference on Artificial Intelligence (AAAI Press, San Jose, CA), AAAI-92, pp 459–465.
6. Davis M, Logemann G, Loveland DW (1962) A machine program for theorem proving. *Commun ACM* 5:394–397.
7. Mézard M, Parisi G, Zecchina R (2002) Analytic and algorithmic solutions of random satisfiability problems. *Science* 297:812–815.
8. Aarts E, Lenstra JK, eds (1997) *Local Search for Combinatorial Optimization* (Wiley, New York).
9. Hoos HH, Stützle T (2005) *Stochastic Local Search: Foundations and Applications* (Elsevier, Amsterdam).
10. Selman B, Kautz H, McAllester D (1997) *Ten Challenges in Propositional Reasoning and Search*, Proceedings of the 15th International Joint Conference on Artificial Intelligence (Elsevier, Amsterdam), IJCAI-97, pp 50–54.
11. Kautz H, Selman B (2007) The state of SAT. *Discrete Appl Math* 155:1514–1524.
12. Kirkpatrick S, Gelatt SD, Jr, Vecchi MP (1983) Optimization by simulated annealing. *Science* 220:671–680.
13. Papadimitriou CH (1991) *On Selecting a Satisfying Truth Assignment*, Proceedings of the 32nd IEEE Symposium on the Foundations of Computer Science (IEEE Computer Society, New York), FOCS-91, pp 163–169.
14. Barthel W, Hartmann AK, Weigt M (2003) Solving satisfiability problems by fluctuations: The dynamics of stochastic local search algorithms. *Phys Rev E* 67:066104.
15. Semerjian G, Monasson R (2003) Relaxation and metastability in a local search procedure for the random satisfiability problem. *Phys Rev E* 67:066103.
16. Selman B, Kautz H, Cohen B (1996) in *Cliques, Coloring, and Satisfiability*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, eds Johnson DS, Trick MA (Am Math Soc, Providence RI), Vol 26, pp 521–532
17. Ardelius J, Aurell E (2006) Behavior of heuristics on large and hard satisfiability problems. *Phys Rev E* 74:037702.
18. Aurell E, Gordon U, Kirkpatrick S (2005) *Comparing Beliefs, Surveys and Random Walks*. Advances in Neural Information Processing Systems (NIPS-04) (MIT Press, Boston), Vol 17, pp 49–56 .
19. Seitz S, Alava M, Orponen P (2005) Focused local search for random 3-satisfiability. *J Stat Mech Theor Exp* 06006.
20. Mertens S, Mézard M, Zecchina R (2006) Threshold values for random *K*-SAT from the cavity method. *Rand Struct Algorithms* 28:340–373.
21. Mézard M, Mora T, Zecchina R (2005) Clustering of solutions in the random satisfiability problem. *Phys Rev Lett* 94:197205.
22. Krzakala F, Montanari A, Ricci-Tersenghi F, Semerjian G, Zdeborova L (2007) Gibbs states and the set of solutions of random constraint satisfaction problems. *Proc Natl Acad Sci USA* 104:10318–10323.
23. Battaglia D, Braunstein A, Chavas J, Zecchina R (2005) Source coding by efficient selection of ground-state clusters. *Phys Rev E* 72:015103.
24. Parisi G (2002) On local equilibrium equations for clustering states. *Technical Report* cs.CC/0212047. Available at http://arxiv.org/abs/cs/0212047.
25. Semerjian G (2008) On the freezing of variables in random constraint satisfaction problems. *J Stat Phys* 130:251–293.
26. Kroc L, Sabharwal A, Selman B (2007) *Survey Propagation Revisited*, Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI-07) (AUAI Press, Arlington, VA), pp 217–226.