

CIRCUMSPECT DESCENT PREVAILS IN SOLVING RANDOM CONSTRAINT SATISFACTION PROBLEMS

MIKKO ALAVA, JOHN ARDELIUS, ERIK AURELL, PETTERI KASKI,
SUPRIYA KRISHNAMURTHY, PEKKA ORPONEN, AND SAKARI SEITZ

ABSTRACT. We study the performance of stochastic local search algorithms for random instances of the K -satisfiability (K -SAT) problem. We introduce a new 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 considers only 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.

1. INTRODUCTION

1.1. **Background.** 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, \dots, 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 [9] constraint satisfaction problem is the K -satisfiability (K -SAT) problem [7], 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; the ground states of the Hamiltonian correspond to the solutions, that is, assignments of values to the variables that satisfy all the clauses (see [14]).

It was first observed in the context of K -SAT, and then in the context of several other CSPs [8], that ensembles of random CSPs have a “phase transition,” a sharp change in the likelihood to be solvable [18]. Empirically, algorithms have been observed to fail or have difficulties in the immediate neighbourhood of such phase transition points, a fact which has given rise to a large literature [8]. 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 [17] motivated by spin glass theory, or variants of stochastic local search techniques [1, 10, 25].

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 neighbourhood 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 [21], 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 [4, 26]. A subsequent influential development occurred with Selman, Kautz and Cohen’s WalkSAT algorithm [24], 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 [2, 3, 23]. In comparison, the satisfiability/unsatisfiability threshold of random 3-satisfiability is believed to be at $\alpha = 4.267$ clauses per variable [15].

1.2. The present work. The present work carries out a first systematic empirical study of random K -SAT for $K = 4$. Our motivation for this study is threefold.

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 [2, 3, 23]. Thus, a similar investigation for higher K is warranted. Here we focus on $K = 4$.

The structure of the space of solutions. Recent rigorous results and non-rigorous 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 [16] 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 a pair of solutions with normalized Hamming distance $0 \leq x \leq 1$. Mézard, Mora and Zecchina [16] show that, for $K \geq 8$, there exists an interval (a, b) , $0 < a < b < 1/2$, such that, with high probability as $N \rightarrow \infty$, 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 $K \geq 8$, 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. [13]) 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.

The structure of the energy landscape. In the context of random K -SAT it is common folklore that SLS algorithms appear to benefit from circumspect descent in energy, that is, from a very conservative policy of lowering the number of clauses not satisfied by the trial configuration. To explore this issue further, we introduce a new SLS algorithm which we call *ChainSAT*. It is based on three ideas: (1) focusing, (2) easing difficult-to-satisfy constraints by so-called *chaining* moves, and (3) *never going upwards in energy*; that is, the number of unsatisfied clauses is a non-increasing 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 a

```

1:  $S$  = random assignment of values to the variables
2: while  $S$  is not a solution do
3:    $C$  = a clause not satisfied by  $S$  selected uniformly at random
4:    $V$  = a variable in  $C$  selected uniformly at random
5:    $\Delta E$  = change in the number of unsatisfied clauses if  $V$  is flipped in  $S$ 
6:   if  $\Delta E \leq 0$  then
7:     flip  $V$  in  $S$ 
8:   else
9:     with probability  $\eta^{\Delta E}$ 
10:    flip  $V$  in  $S$ 
11:   end with
12: end if
13: end while

```

FIGURE 1. The Focused Metropolis Search algorithm [23].

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 [23]. This observation further supports the position that random K -SAT for small K may be more elementary than has been previously believed.

1.3. Organization of the paper. Section 2 documents our experiments with the FMS algorithm on random K -SAT for $K = 4$. Section 3 contains an empirical investigation of x -satisfiability in random K -SAT for $K = 4$ using the FMS algorithm. Section 4 introduces the ChainSAT algorithm and studies its performance on random K -SAT for $K = 4, 5, 6$. Section 5 presents a few concluding remarks.

2. EXPERIMENTS WITH FOCUSED METROPOLIS SEARCH

The Focused Metropolis Search (FMS) algorithm [23] is given in pseudocode in Figure 1. This section documents our experiments aimed at charting the empirical linear-time region of FMS on random K -SAT for $K = 4$.

2.1. Selecting the temperature parameter. For $K = 3$ it has already been established that the FMS algorithm has an “operating window” in terms of the adjustable “temperature” parameter η [23]. 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 .

Figure 2 shows a typical result of the optimization of the temperature parameter η for random K -SAT with $K = 4$. Two quantities are plotted, the fraction of instances solved (within a threshold number of flips per variable), and, when all instances are solved, the corresponding average solution time.

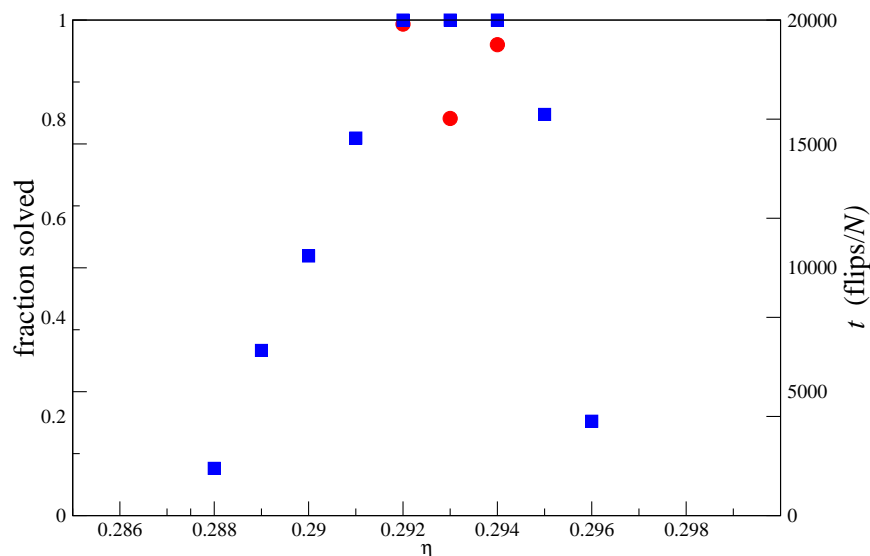


FIGURE 2. Optimizing the temperature parameter η for the FMS algorithm on instances of random K -SAT at $K = 4$ and $\alpha = 9.6$. Displayed on the horizontal axis is the temperature parameter η . Plotted on the vertical axis is the fraction of 21 random instances solved within $60000 \times N$ flips at $N = 100000$. In the case all 21 instances are solved, also plotted is the average solution time (in flips/ N). The optimum is at $\eta = 0.293$. Note the narrowness of the operating window in terms of η .

2.2. The empirical linear-time regime of FMS. It is evident from Figure 2 that for $K = 4$ and $\alpha = 9.6$ the operating window of FMS is already very narrow; thus it is striking that the empirical performance of FMS is almost surely linear in N within the window.

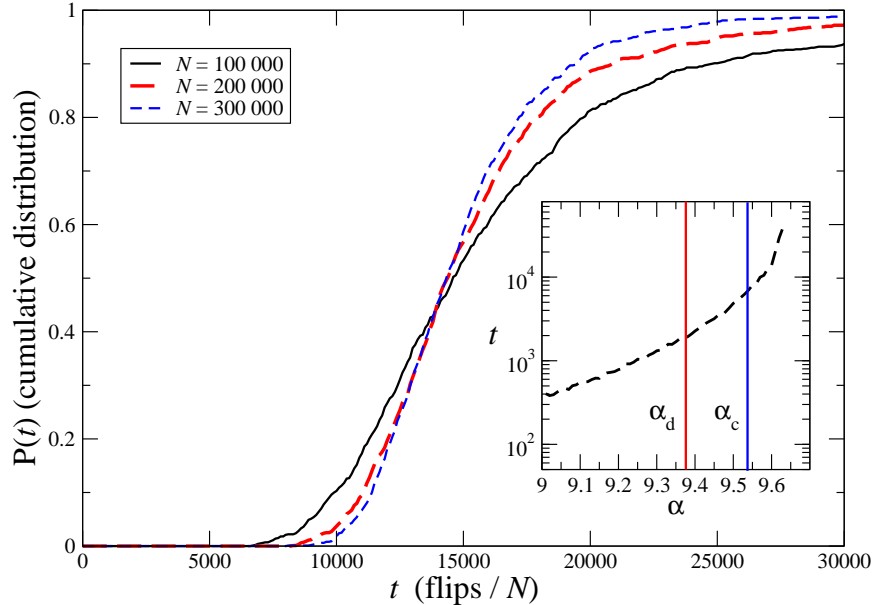


FIGURE 3. Cumulative distributions of solution times normalized by the number of variables N for the Focused Metropolis Search algorithm [23] on instances of random K -satisfiability at $K = 4$ and $\alpha = 9.6$. The vertical axis indicates the fraction of 1001 random instances solved within a given 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 $\eta = 0.293$.

In Figure 3 we present empirical evidence that FMS almost surely runs in time linear in N for instances of random K -satisfiability with $K = 4$. The fact that the curves get steeper with increasing N implies concentration of solution times, or that above-average and below-average solution times get rarer with 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 [13] spin-glass theoretic “dynamical” and “condensation” transition points.

3. 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” transition points [13]. In particular, Figure 4 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. for a balanced version of $K = 5$ [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 (see [23], [22], and Section 4.3)—all the solutions found in our experiments were completely white, that is, they do not have locally frozen variables. One can of course imagine that a “typical solution” is not white, under the circumstances examined here, but as noted there is no evidence of the existence of such.

Figure 5 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 .

Figure 6 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.

4. EXPERIMENTS WITH CHAINSAT

A new heuristic which never moves up in energy is here shown to solve random K -satisfiability problems almost surely in time linear in N , for $K = 4, 5, 6$.

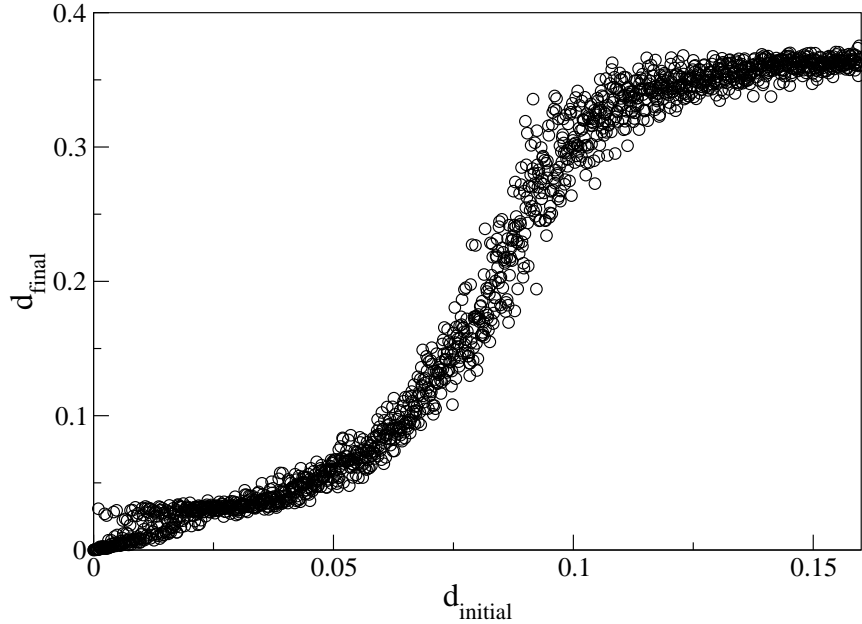


FIGURE 4. Investigation of x -satisfiability using FMS initialized with a random configuration at a given Hamming distance from a reference solution. One reference solution and one instance of random K -SAT at $K = 4$, $\alpha = 9.6$, and $N = 200000$. The horizontal axis displays the normalized Hamming 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 of the 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 $\eta = 0.293$.

4.1. The ChainSAT algorithm. Our new heuristic, ChainSAT, is given in pseudocode in Figure 7. The algorithm (a) never increases the energy of the current configuration S ; and (b) exercises circumspection in decreasing the energy. In particular, moves that decrease the energy are taken only sporadically compared with equi-energetic moves and chaining moves. The latter are designed to alleviate critically satisfied constraints by proceeding in “chains” of variable-clause-variable until

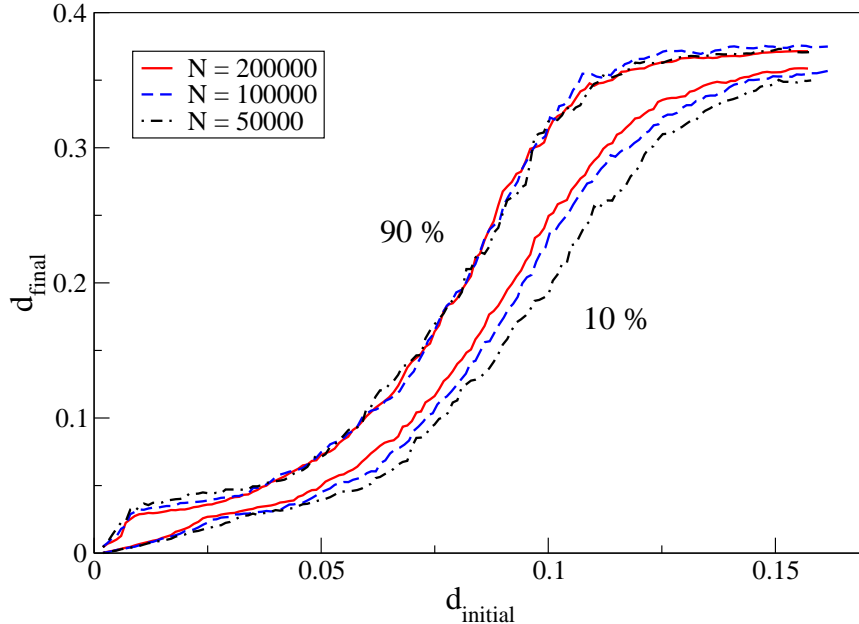


FIGURE 5. Scaling of x -satisfiability data obtained using FMS on random K -SAT with increasing N . The parameters $K = 4$, $\alpha = 9.6$, and $\eta = 0.293$ are fixed. The plotted 10- and 90-percentile curves are calculated from five random instances and reference solutions for each $N = 50000, 100000, 200000$, with a moving window size of 0.004 in the horizontal axis. The distances appear to converge close to the 90%-curves.

a variable is found which can be flipped without increase in energy. Focusing is employed for the non-chaining 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. We omit data related to the optimization of these parameters since the procedure is simply an empirical (vary the parameters, check outcome), similar to one documented for the FMS algorithm in Figure 2.

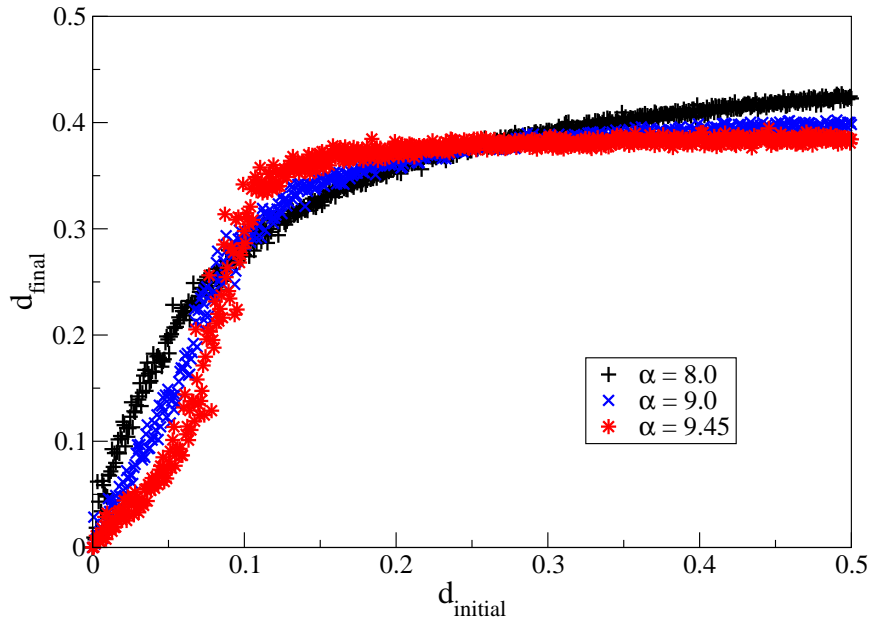


FIGURE 6. Scaling of x -satisfiability data obtained using FMS on random K -SAT with increasing α . The parameters $K = 4$, $N = 100000$, and $\eta = 0.293$ are fixed. One random instance and one reference solution for each $\alpha = 8.0, 9.0, 9.45$; see Figure 4 for $\alpha = 9.6$. The value $\alpha = 9.45$ is between the predicted locations of the dynamical and the condensation transition points [13]. No clear gap in distances is discernible in any of the cases.

4.2. ChainSAT performance. In Figure 8 we present empirical evidence that ChainSAT almost surely runs in time linear in N for random K -satisfiability problems with $K = 4, 5, 6$. The fact that the curves get steeper with increasing N implies concentration of solution times, or that above-average and below-average solution times get rarer with N .

Since 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.

```

1:  $S$  = random assignment of values to the variables
2: chaining = FALSE
3: while  $S$  is not a solution do
4:   if not chaining then
5:      $C$  = a clause not satisfied by  $S$  selected uniformly at random
6:      $V$  = a variable in  $C$  selected u.a.r.
8:   end if
9:    $\Delta E$  = change in the number of unsatisfied clauses if  $V$  is flipped in  $S$ 
10:  chaining = FALSE
11:  if  $\Delta E = 0$  then
12:    flip  $V$  in  $S$ 
13:  else if  $\Delta E < 0$ 
14:    with probability  $p_1$ 
15:    flip  $V$  in  $S$ 
16:    end with
17:  else
18:    with probability  $1 - p_2$ 
19:     $C$  = a clause satisfied only by  $V$  selected u.a.r.
20:     $V'$  = a variable in  $C$  other than  $V$  selected u.a.r.
21:     $V = V'$ 
22:    chaining = TRUE
23:    end with
24:  end if
25: end while

```

FIGURE 7. The ChainSAT algorithm.

4.3. **Whiteness.** To provide a further empirical analysis of ChainSAT, we next present Figure 9. 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 $l_{\text{chain}} = f/m - 1$, 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 [22], described in pseudocode in Figure 10, that 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

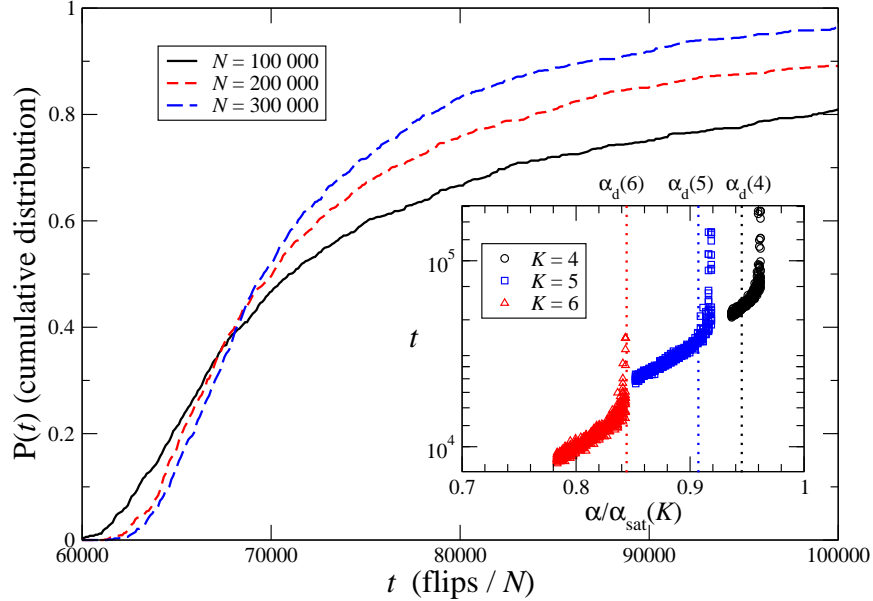


FIGURE 8. Cumulative distributions of solution times normalized by number of variables N for the ChainSAT algorithm on random K -satisfiability instances at $K = 4$ and $\alpha = 9.55$. The vertical axis indicates the fraction of 1001 random input instances solved within a given running time, measured 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 $\alpha_{\text{sat}}(K)$, which has the empirical values $\alpha_{\text{sat}}(4) = 9.931$, $\alpha_{\text{sat}}(5) = 21.117$, and $\alpha_{\text{sat}}(6) = 43.37$ [15]. The parameters of ChainSAT have been chosen to be small enough to work at least up to the predicted “dynamical transition” [13]: we have set $p_1 = p_2 = 0.0001$ ($K = 4$), 0.0002 ($K = 5$), and 0.0005 ($K = 6$).

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 [23] for an empirical discussion of AWD 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.

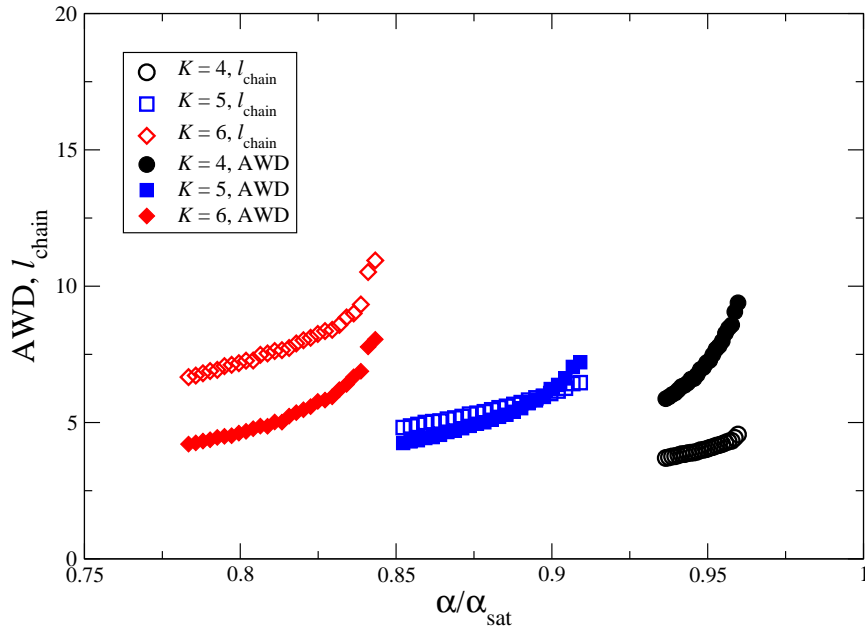


FIGURE 9. The average chain length in ChainSAT and the average whiteness depth of the solutions found in random K -SAT for $K = 4, 5, 6$. Each plotted value is the average over 21 random instances. The values of $\alpha(K)$ in the horizontal axis have been normalized with $\alpha_{\text{sat}}(K)$, which has the empirical values $\alpha_{\text{sat}}(4) = 9.931$, $\alpha_{\text{sat}}(5) = 21.117$, and $\alpha_{\text{sat}}(6) = 43.37$ [15]. The ChainSAT parameters are set to $p_1 = p_2 = 0.0001$ ($K = 4$), 0.0002 ($K = 5$), and 0.0005 ($K = 6$).

Based on Figure 9 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 $\text{AWD}/l_{\text{chain}}$ increases with α .

5. CONCLUDING REMARKS

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

```

1: initially all clauses and variables are unmarked (non-white)
2: mark (whiten) every clause that is unsatisfied
3: mark (whiten) every clause that has more than one true literal
4:  $D = 0$ 
5: repeat
6:     mark (whiten) any unmarked variables that appear as satisfying
       literals only in marked clauses
7:     if all the variables are marked then
8:         declare that  $S$  is completely white
9:         halt
10:    end if
11:    if no new variables were marked in this iteration then
12:        declare that  $S$  has a core
13:        halt
14:    end if
15:    mark (whiten) any unmarked clauses that contain at least
       one marked variable
16:     $D = D + 1$ 
17: end repeat

```

FIGURE 10. The whitening algorithm for a configuration S .

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 towards 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 $K \geq 8$ [16] 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 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?

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.

Acknowledgements: This work was supported by the Integrated Project EVERGROW of the European Union (J.A and S.K), by the Swedish Science Council through Linnaeus Centre ACCESS (E.A.), and by the Academy of Finland under Grant 117499 (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 SICS for the use of a little over 6 years of CPU time.

REFERENCES

- [1] E. Aarts and J. K. Lenstra (Eds.), *Local Search for Combinatorial Optimization*. J. Wiley & Sons, New York NY, 1997.
- [2] J. Ardelius and E. Aurell, Behavior of heuristics on large and hard satisfiability problems. *Phys. Rev. E* 74 (2006) 037702.
- [3] E. Aurell, U. Gordon, and S. Kirkpatrick. Comparing beliefs, surveys and random walks. *Advances in Neural Information Processing Systems* 17 (NIPS-04), 49–56. The MIT Press, 2005.
- [4] W. Barthel, A. K. Hartmann, and M. Weigt, Solving satisfiability problems by fluctuations: The dynamics of stochastic local search algorithms. *Phys. Rev. E* 67 (2003) 066104.
- [5] D. Battaglia, A. Braunstein, J. Chavas, and R. Zecchina, Source coding by efficient selection of ground-state clusters, *Phys. Rev. E* 72 (2005) 015103.
- [6] M. Davis, G. Logemann and D. W. Loveland. A machine program for theorem proving. *Communications of the ACM* 5 (1962) 394–397.
- [7] D. Du, J. Gu, P. Pardalos (Eds.), *Satisfiability Problem: Theory and Applications*. DIMACS Series in Discr. Math. and Theoret. Comput. Sci. 35, American Math. Soc., Providence RI, 1997.
- [8] O. Dubois, R. Monasson, B. Selman and R. Zecchina (Eds.). Special Issue on NP-Hardness and Phase Transitions. *Theoret. Comput. Sci.* 265 (2001).
- [9] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., San Francisco CA, 1979.

- [10] H. H. Hoos and T. Stützle, *Stochastic Local Search: Foundations and Applications*. Morgan Kaufmann (Elsevier), 2005.
- [11] H. Kautz and B. Selman. The state of SAT. *Discrete Appl. Math.* 155 (2007) 1514–1524.
- [12] S. Kirkpatrick, C. D. Gelatt, Jr., and M. P. Vecchi, Optimization by simulated annealing, *Science* 220 (1983) 671–680.
- [13] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborova, Gibbs states and the set of solutions of random constraint satisfaction problems. *Proc. Nat. Acad. Sci. USA* 104 (2007) 10318–10323.
- [14] O. C. Martin, R. Monasson, and R. Zecchina, Statistical mechanics methods and phase transitions in optimization problems, *Theoret. Comput. Sci.* 265 (2001) 3–67.
- [15] S. Mertens, M. Mézard, and R. Zecchina, Threshold values for random K -SAT from the cavity method, *Random Structures Algorithms* 28 (2006) 340–373.
- [16] M. Mézard, T. Mora, and R. Zecchina, Clustering of solutions in the random satisfiability problem, *Phys. Rev. Lett.* 94 (2005) 197205.
- [17] M. Mézard, G. Parisi, and R. Zecchina, Analytic and algorithmic solutions of random satisfiability problems. *Science* 297 (2002) 812–815.
- [18] D. Mitchell, B. Selman, and H. Levesque, Hard and easy distributions of SAT problems. *Proc. 10th Natl. Conf. on Artificial Intelligence (AAAI-92)*, 459–465. AAAI Press, San Jose CA, 1992.
- [19] R. Monasson and R. Zecchina, Statistical mechanics of the random K -SAT problem. *Phys. Rev. E* 56 (1997) 1357–1361.
- [20] A. Montanari, G. Parisi, and F. Ricci-Tersenghi, Instability of one-step replica-symmetry-broken phase in satisfiability problems, *J. Phys. A* 37 (2004) 2073–2091.
- [21] C.H. Papadimitriou, On selecting a satisfying truth assignment. *Proc. 32nd IEEE Symposium on the Foundations of Computer Science (FOCS-91)*, 163–169. IEEE Computer Society, New York NY, 1991.
- [22] G. Parisi, On local equilibrium equations for clustering states. Technical report cs.CC/0212047, arXiv.org (Feb 2002).
- [23] S. Seitz, M. Alava, and P. Orponen, Focused local search for random 3-satisfiability, *J. Stat. Mech. Theory Exp.* (2005) 06006.
- [24] B. Selman, H. Kautz, and B. Cohen, Local search strategies for satisfiability testing. In: D. S. Johnson and M. A. Trick (Eds.), *Cliques, Coloring, and Satisfiability*, 521–532. DIMACS Series in Discr. Math. and Theoret. Comput. Sci. 26, American Math. Soc., Providence RI, 1996.
- [25] B. Selman, H. Kautz, and D. McAllester. Ten challenges in propositional reasoning and search. *Proceedings, 15th International Joint Conference on Artificial Intelligence (IJCAI-97)*, 50–54. Morgan Kaufmann, 1997.
- [26] G. Semerjian and R. Monasson, Relaxation and metastability in a local search procedure for the random satisfiability problem. *Phys. Rev. E* 67 (2003) 066103.

LABORATORY OF PHYSICS, P.O. BOX 1100, FI-02015 HELSINKI UNIVERSITY OF TECHNOLOGY, FINLAND

E-mail address: mikko.alava@tkk.fi

SICS SWEDISH INSTITUTE OF COMPUTER SCIENCE AB, SE-164 29 KISTA, SWEDEN

DEPARTMENT OF COMPUTATIONAL BIOLOGY, ALBANOVA UNIVERSITY CENTRE, SE-106 91 STOCKHOLM, SWEDEN

DEPARTMENT OF COMPUTATIONAL BIOLOGY, ALBANOVA UNIVERSITY CENTRE, SE-106 91 STOCKHOLM, SWEDEN

ACCESS 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 UNIVERSITY OF HELSINKI, FINLAND

E-mail address: `petteri.kaski@cs.helsinki.fi`

SICS SWEDISH INSTITUTE OF COMPUTER SCIENCE AB, SE-164 29 KISTA, SWEDEN

SCHOOL OF INFORMATION AND COMMUNICATION TECHNOLOGY, KTH-ROYAL INSTITUTE OF TECHNOLOGY, SE-164 40 KISTA, SWEDEN

LABORATORY FOR THEORETICAL COMPUTER SCIENCE, P.O. Box 5400, FI-02015 HELSINKI UNIVERSITY OF TECHNOLOGY, FINLAND

E-mail address: `pekka.orponen@tkk.fi`

LABORATORY FOR THEORETICAL COMPUTER SCIENCE, P.O. Box 5400, FI-02015 HELSINKI UNIVERSITY OF TECHNOLOGY, FINLAND

E-mail address: `sakari.seitz@tkk.fi`