Threshold Behaviour of WalkSAT and Focused Metropolis Search on Random 3-Satisfiability*

Sakari Seitz[1,2], Mikko Alava[2], and Pekka Orponen[1]

[1] Laboratory for Theoret. Computer Science, P.O.B. 5400
[2] Laboratory of Physics, P.O.B. 1100
Helsinki University of Technology
FI-02015 TKK, Finland
E-mail: `firstname.lastname@tkk.fi`

Abstract. An important heuristic in local search algorithms for Satisfiability is *focusing*, i.e. restricting the selection of flipped variables to those appearing in presently unsatisfied clauses. We consider the behaviour on large randomly generated 3-SAT instances of two focused solution methods: the well-known WalkSAT algorithm, and the straightforward algorithm obtained by imposing the focusing constraint on basic Metropolis dynamics. We observe that both WalkSAT and our Focused Metropolis Search method are quite sensitive to the proper choice of their "noise" and "temperature" parameters, and attempt to estimate ideal values for these, such that the algorithms exhibit generically linear solution times as close to the satisfiability transition threshold $\alpha_c \approx 4.267$ as possible. With an appropriate choice of parameters, the linear time regime of both algorithms seems to extend well into clauses-to-variables ratios $\alpha > 4.2$, which is much further than has generally been assumed in the literature.

1 Introduction

An instance of the 3-satisfiability (3-SAT) problem is a formula consisting of M clauses, each of which is a set of three literals, i.e. Boolean variables or their negations. The goal is to find a solution consisting of a satisfying truth assignment to the N variables, such that each clause contains at least one literal evaluating to 'true', provided such an assignment exists. In a random 3-SAT instance, the literals comprising each clause are chosen uniformly at random.

It was observed in [15] that random 3-SAT instances change from being generically satisfiable to being generically unsatisfiable when the clauses-to-variables ratio $\alpha = M/N$ exceeds a critical threshold α_c. Current numerical estimates [4] suggest that this satisfiability threshold is located approximately at $\alpha_c \approx 4.267$. For a general introduction to aspects of the satisfiability problem see [6].

Two questions are often asked in this context: how to solve random 3-SAT instances effectively close to the threshold α_c, and what can be said of the point

* Research supported by the Academy of Finland, by grants 206235 (S. Seitz) and 204156 (P. Orponen), and by the Center of Excellence program (M. Alava).

at which various types of algorithms begin to falter. Recently progress has been made by the *survey propagation* method [4, 5, 14] that essentially iterates a guess about the state of each variable, in the course of fixing an ever larger proportion of the variables to their "best guess" values. In this paper we argue that also simple local search methods, i.e. algorithms that try to find a solution by flipping the value of one variable at a time, can achieve comparable performance levels on random 3-SAT instances.[1]

When the variables to be flipped are chosen only from the unsatisfied (unsat) clauses, a local search algorithm is called *focusing.* A well-known example of a focused local search algorithm for 3-SAT is WalkSAT [23], which makes greedy and random flips with a fixed probability. Many variants of this and other local search algorithms with different heuristics have been developed; for a general overview of the techniques see [1]. We shall contrast the WalkSAT algorithm with a focused variant of the standard Metropolis dynamics [12] of spin systems, which is also the base for the well-known simulated annealing optimisation method. When applied to the 3-SAT problem, we call this dynamics the Focused Metropolis Search (FMS) algorithm.

The space of solutions to 3-SAT instances slightly below α_c is known to develop structure. Physics-based methods from mean-field -like spin glasses imply through replica methods that the solutions become "clustered", with a threshold of $\alpha \approx 3.92$ [14]. Clustering implies that solutions belonging to the same cluster are close to each other in terms of, e.g., Hamming distance. A possible consequence of this is the existence of cluster "backbones", which means that in a cluster a certain fraction of the variables is fixed, while the others can be varied subject to some constraints. The stability of the replica ansatz becomes crucial for higher α, and it has been suggested that this might become important at $\alpha \approx 4.15$ [16]. The energy landscape in which the local algorithms move is however a finite-temperature version. The question is how close to α_c one can get by moving around by local moves (spin flips) and focusing on the unsat clauses.

Our results concern the optimality of this strategy for various algorithms and clauses-to-variables ratios α. First, we demonstrate that WalkSAT works in the "critical" region up to $\alpha > 4.2$ with the optimal noise parameter $p \approx 0.57$. Then we concentrate on the numerical performance of the FMS method. In this case, the solution time is found to be linear in N within a parameter window $\eta_{min} \leq \eta \leq \eta_{max}$, where η is essentially the Metropolis temperature of the FMS dynamics. More precisely, we observe that within this window, the median and all other quantiles of the solution times normalised by N seem to approach a constant value as N increases. A stronger condition would be that the distributions of the solution times be concentrated in the sense that the ratio of the standard deviation and the average solution time decreases with increasing N. While numerical studies of the distributions do not indicate heavy tails that would contradict this condition, we can of course not at present prove that this

[1] Anecdotal evidence suggests that local search methods may also be more robust than survey propagation on structured SAT instances, but we have at present no systematical data to support this conjecture.

is rigorously true. The existence of the η window implies that for too large η the algorithm becomes entropic, while for too small η it is too greedy, leading to freezing of degrees of freedom (variables) in that case.

The *optimal* $\eta = \eta_{opt}(\alpha)$, i.e. that η for which the solution time median is lowest, seems to increase with increasing α. We have tried to extrapolate this behaviour towards α_c and consider it possible that the algorithm works, in the median sense, linearly in N all the way up to α_c. This is in contradiction with some earlier conjectures. All in all we postulate a phase diagram for the FMS algorithm based on two phase boundaries following from too little or too much greed. Of this picture, one particular phase space point has been considered before [3, 25], since it happens to be the case that for $\eta = 1$ the FMS coincides with the so called Random Walk method [17], which is known to suffer from metastability at $\alpha \approx 2.7$ [3, 25].

The structure of this paper is as follows: In Section 2 we present the WalkSAT algorithm, list some of the known results concerning the random 3-SAT problem, and report on our experiments on the behaviour of WalkSAT for varying p and α values. In Section 3 we outline the FMS algorithm, report on the corresponding numerical simulations with it, and present some analysis of what the data indicates about this algorithm's threshold behaviour. Section 4 summarises our results. An extended version of this paper, available as a technical report [22], contains some further experiments covering also the so called Focused Record-to-Record Travel (FRRT) algorithm [7, 21], as well as a discussion on how the notion of *whitening* [18] might shed some light on the dynamics of focused local search.

2 Local Search for Satisfiability

It is natural to view the satisfiability problem as a combinatorial optimisation task, where the goal for a given formula F is to minimise the objective function $E = E_F(s) =$ the number of unsatisfied clauses in formula F under truth assignment s. The use of local search heuristics in this context was promoted e.g. by Selman et al. in [24] and by Gu in [8]. Viewed as a spin glass model, E can be taken to be the energy of the system.

Selman et al. introduced in [24] the simple greedy GSAT algorithm, whereby at each step the variable to be flipped is determined by which flip yields the largest decrease in E, or if no decrease is possible, then smallest increase. This method was improved in [23] by augmenting the simple greedy steps with an adjustable fraction p of pure random walk moves, leading to the algorithm NoisyGSAT.

In a different line of work, Papadimitriou [17] introduced the very important idea of *focusing* the search to consider at each step only those variables that appear in the presently *unsatisfied* clauses. Applying this modification to the NoisyGSAT method leads to the WalkSAT [23] algorithm (Figure 1), which is arguably the currently most popular local search method for satisfiability.

```
WalkSAT(F,p):
  s = random initial truth assignment;
  while flips < max_flips do
    if s satisfies F then output s & halt, else:
    - pick a random unsatisfied clause C in F;
    - if some variables in C can be flipped without
      breaking any presently satisfied clauses,
      then pick one such variable x at random; else:
    - with probability p, pick a variable x
      in C at random;
    - with probability (1-p), pick some x in C
      that breaks a minimal number of presently
      satisfied clauses;
    - flip x.
```

Fig. 1. The WalkSAT algorithm.

In [23], Selman et al. presented some comparisons among their new NoisyGSAT and WalkSAT algorithms, together with some other methods. These experiments were based on a somewhat unsystematic set of benchmark formulas with $N \leq 2000$ at $\alpha \approx \alpha_c$, but nevertheless illustrated the significance of the focusing idea, since in the results reported, WalkSAT outperformed NoisyGSAT by several orders of magnitude.

More recently, Barthel et al. [3] performed systematic numerical experiments with Papadimitriou's original Random Walk method at $N = 50,000$, $\alpha = 2.0 \ldots 4.0$. They also gave an analytical explanation for a transition in the dynamics of this algorithm at $\alpha_{\mathrm{dyn}} \approx 2.7$, already observed by Parkes [20]. When $\alpha < \alpha_{\mathrm{dyn}}$, satisfying assignments are generically found in time that is linear in the number of variables, whereas when $\alpha > \alpha_{\mathrm{dyn}}$ exponential time is required. (Similar results were obtained by Semerjian and Monasson in [25], though with smaller experiments ($N = 500$).) The reason for this dynamical threshold phenomenon seems to be that at $\alpha > \alpha_{\mathrm{dyn}}$ the search equilibrates at a nonzero energy level, and can only escape to a ground state through a large enough random fluctuation. A rate-equation analysis of the method [3] yields a very well matching approximation of $\alpha_{\mathrm{dyn}} \approx 2.71$.[2] See also [26] for further analyses of the Random Walk method on random 3-SAT.

WalkSAT is more powerful than the simple Random Walk, because in it focusing is accompanied by other heuristics. However, it is known that the behaviour of WalkSAT is quite sensitive to the choice of the noise parameter p. E.g. Parkes [19] experimented with the algorithm using a noise value $p = 0.3$ and concluded that with this setting the algorithm works in linear time at least up to $\alpha = 3.8$. Even this result is not the best possible, since it has been estimated [9, 10, 27] that for random 3-SAT close to the satisfiability transition the optimal

[2] Our numerical experiments with the Random Walk algorithm suggest that its dynamical threshold is actually somewhat lower, $\alpha_{\mathrm{dyn}} \approx 2.67$.

noise setting for WalkSAT is $p \approx 0.55$. (Actually our numerical experiments, reported below, suggest that the ideal value is closer to $p \approx 0.57$.)

These positive results notwithstanding, it has been generally conjectured (e.g. in [4]) that no local search algorithm can work in linear time beyond the clustering transition at $\alpha_d \approx 3.92$. In a series of recent experiments, however, Aurell et al. [2] concluded that with a proper choice of parameters, the median solution time of WalkSAT remains linear in N up to at least $\alpha = 4.15$, the onset of 1-RSB symmetry breaking. Our experiments, reported below, indicate that the median time in fact remains linear even beyond that, in line with our previous results [21].

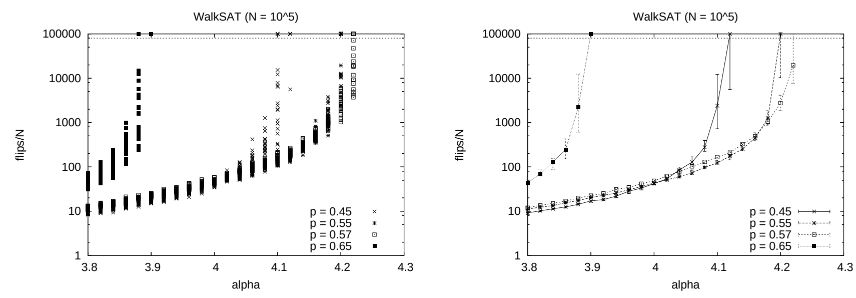

(a) Complete data

(b) Medians and quartiles

Fig. 2. Normalised solution times for WalkSAT, $\alpha = 3.8 \ldots 4.3$.

Figure 2 illustrates our experiments with the WalkSAT algorithm[3] on randomly generated formulas of size $N = 10^5$, various values of the noise parameter p, and values of α starting from 3.8 and increasing at increments of 0.2 up to 4.22. For each (p, α)-combination, 21 formulas were generated, and for each of these the algorithm was run until either a satisfying solution was found or a time limit of $80000 \times N$ flips was exceeded. Figure 2(a) shows the solution times t_{sol}, measured in number of flips normalised by N, for each generated formula. Figure 2(b) gives the medians and quartiles of the data for each value of α.

As can be seen from the figures, for value $p = 0.45$ of the noise parameter WalkSAT finds satisfying assignments in roughly time linear in N, with the coefficient of linearity increasing gradually with increasing α, up to approximately $\alpha = 4.1$ beyond which the distribution of solution times for the algorithm diverges. For $p = 0.55$, this linear regime extends further, up to at least $\alpha = 4.18$, but for $p = 0.65$ it seems to end already before $\alpha = 3.9$. For the best value of

[3] Version 43, downloaded from the Walksat Home Page at `http://www.cs.washington.edu/homes/kautz/walksat/`, with its default heuristics.

p we have been able to experimentally determine, $p = 0.57$, the linear regime seems to extend up to even beyond $\alpha = 4.2$.

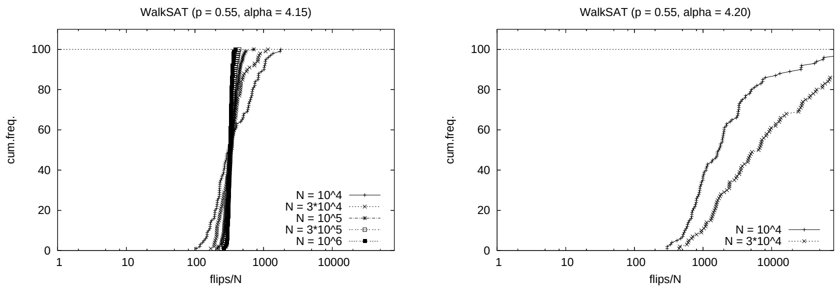

(a) $\alpha = 4.15$ (b) $\alpha = 4.20$

Fig. 3. Cumulative solution time distributions for WalkSAT with $p = 0.55$.

To investigate the convergence of the solution time distributions, we tested the WalkSAT algorithm with $p = 0.55$ at $\alpha = 4.15$ and $\alpha = 4.20$, in both cases with randomly generated sets of 100 formulas of sizes $N = 10^4, 3\times10^4, 10^5, 3\times10^5$ and 10^6. Figure 3 shows the cumulative distributions of the solution times normalised by N achieved in these tests. As can be seen, for $\alpha = 4.15$ the distributions are well-defined, with normalised medians and all other quantiles converging to a finite value for increasing N. However, for $\alpha = 4.20$, the distributions seem to diverge, with median values increasing with increasing N.[4]

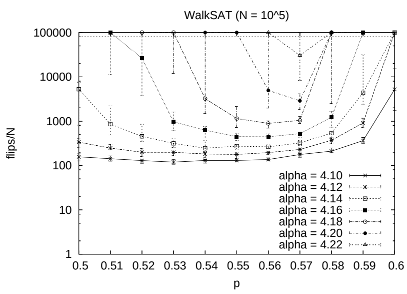

Fig. 4. Normalised solution times for WalkSAT with $\alpha = 4.10 \ldots 4.22$, $p = 0.50 \ldots 0.60$.

[4] For $\alpha = 4.20$, the tests for $N = 10^5$ and larger were not completed, because the solution times consistently overran the time limit of $80000\times N$ flips, and consequently the test runs were exceedingly long yet uninformative.

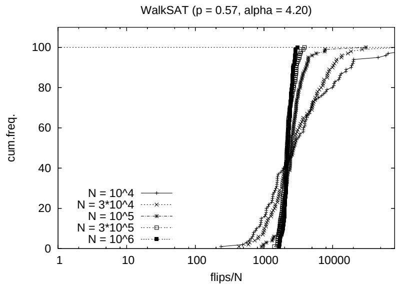

Fig. 5. Cumulative solution time distributions for WalkSAT with $p = 0.57$, $\alpha = 4.20$.

In order to estimate the optimal value of the WalkSAT noise parameter, i.e. that value of p for which the linear time regime extends to the biggest value of α, we generated test sets consisting of 21 formulas, each of size $N = 10^5$, for α values ranging from 4.10 to 4.22 at increments of 0.02, and for each α for p values ranging from 0.50 to 0.60 at increments of 0.01. WalkSAT was then run on each resulting (α, p) test set; the medians and quartiles of the observed solution time distributions are shown in Figure 4. The data suggest that the optimal value of the noise parameter is approximately $p = 0.57$. Figure 5 shows the empirical solution time distributions at $\alpha = 4.20$ for $p = 0.57$. In contrast to Figure 3(b), now the quantiles seem again to converge to a finite value for increasing N, albeit more slowly than in the case of the simpler $\alpha = 4.15$ formulas presented in Figure 3(a).

3 Focused Metropolis Search

From an analytical point of view, the WalkSAT algorithm is rather complicated, with its interleaved greedy and randomised moves. Thus, it is of interest to investigate the behaviour of the simpler algorithm obtained by imposing the focusing heuristic on a basic Metropolis dynamics [12].

```
FMS(F,eta):
  s = random initial truth assignment;
  while flips < max_flips do
    if s satisfies F then output s & halt, else:
      pick a random unsatisfied clause C in F;
      pick a variable x in C at random;
      let x' = flip(x), s' = s[x'/x];
      if E(s') <= E(s) then flip x, else:
         flip x with prob. eta^(E(s')-E(s)).
```

Fig. 6. The Focused Metropolis Search algorithm.

We call the resulting algorithm, outlined in Figure 6, the *Focused Metropolis Search* (FMS) method. The algorithm is parameterised by a number η, $0 \leq \eta \leq 1$, which determines the probability of accepting a candidate variable flip that would lead to a unit increase in the objective function E. (Thus in customary Metropolis dynamics terms, $\eta = e^{1/T}$, where T is the chosen computational temperature. Note, however, that detailed balance does not hold with focusing.)

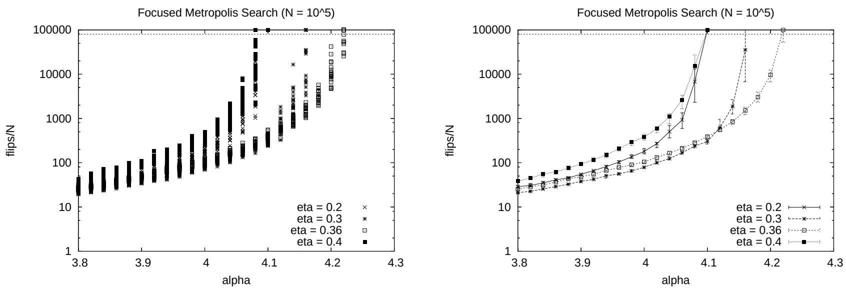

(a) Complete data (b) Medians and quartiles

Fig. 7. Normalised solution times for FMS, $\alpha = 3.8 \ldots 4.3$.

We repeated the data collection procedure of Figure 2 for the FMS algorithm with various parameter values. The results for $\eta = 0.2, 0.3, 0.4$ and $\eta = 0.36$ (the best value we were able to find) are shown in Figure 7; note that also rejected flips are here included in the flip counts. As can be seen, the behaviour of the algorithm is qualitatively quite similar to WalkSAT. For parameter value $\eta = 0.2$, the linear time regime seems to extend up to roughly $\alpha = 4.06$, for $\eta = 0.3$ up to at least $\alpha = 4.14$, and for $\eta = 0.36$ even beyond $\alpha = 4.20$; however for $\eta = 0.4$ again only up to maybe $\alpha = 4.08$. To test the convergence of distributions, we determined the empirical cumulative distributions of FMS solution times for $\eta = 0.3$ at $\alpha = 4.0$ and $\alpha = 4.1$, in a similar manner as in Figure 3. The results are shown in Figure 8.

To determine the optimal value of the η parameter we proceeded as in Figure 4, mapping out systematically the solution time distributions of the FMS algorithm for α increasing from 4.10 to 4.22 and η ranging from 0.28 to 0.38. The results, shown in Figure 9, suggest that the optimal value of the parameter is approximately $\eta = 0.36$.

In order to investigate the algorithm's behaviour at this extreme of its parameter range, we determined the empirical cumulative distributions of the FMS solution times for $\eta = 0.36$ at $\alpha = 4.20$. The results, shown in Figure 10, suggest that even for this high value of α, the FMS solution times are linear in N, with all quantiles converging to a finite value as N increases.

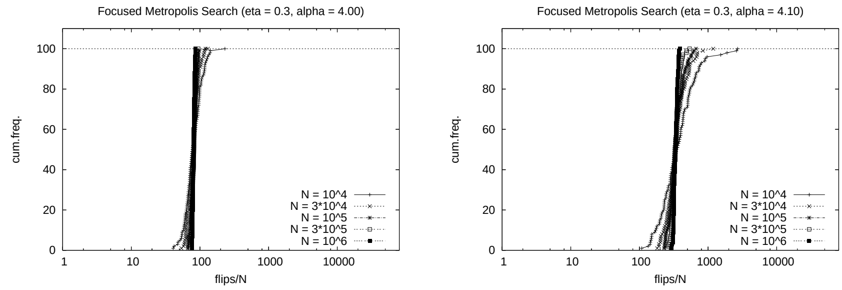

(a) $\alpha = 4.00$ (b) $\alpha = 4.10$

Fig. 8. Cumulative solution time distributions for FMS with $\eta = 0.3$.

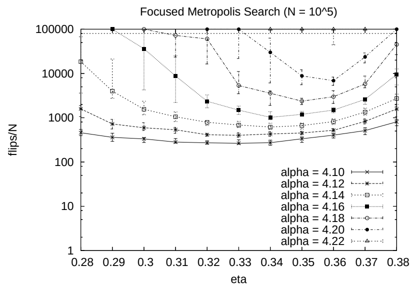

Fig. 9. Normalised solution times for FMS with $\eta = 0.28 \dots 0.38$, $\alpha = 4.10 \dots 4.22$.

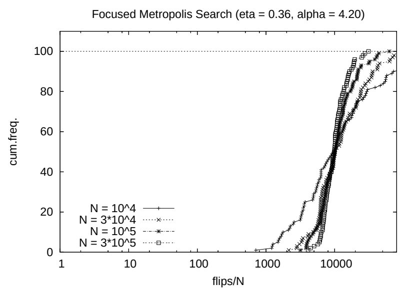

Fig. 10. Cumulative solution time distributions for FMS with $\eta = 0.36$, $\alpha = 4.20$.

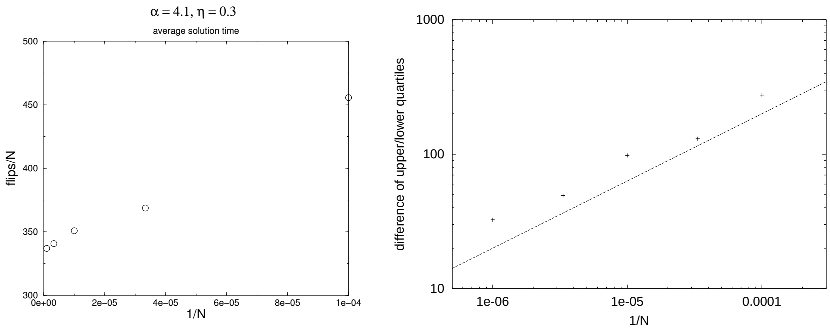

Fig. 11. (a) The N-dependence of the average solution time for $\alpha = 4.1$ and $\eta = 0.3$. (b) The difference of the upper and lower quartiles in t_{sol} vs. N for the same parameters. The line shows $1/\sqrt{N}$-behaviour as a guide for the eye.

The linearity of FMS can fail due to the formation of heavy tails. This, with a given α and a not too optimal, large η would imply that the solution time t_{sol} has at least a divergent mean (first moment) and perhaps a divergent median as well. This can be deliberated upon by considering the "scaling ansatz"

$$P(t_{sol}) \sim (t_{sol})^{-a} f(t_{sol}/N^b) \qquad (1)$$

where $f(x) = \text{const}$ for x small, and $f \to 0$ rapidly for $x \geq 1$. This simply states that for a fixed N there has to be a maximal solution time (even exponentially rare) since the FMS is generically "ergodic" or able to get out of local minima. The condition that $\langle t_{sol} \rangle \sim N^{b(2-a)}$ be divergent with N would then set a relation for the exponents a and b. Our experiments in the linear regime have not yielded any evidence for a distribution following such an ansatz, and moreover we have not explored systematically the non-linear parameter region, where such behaviour might make considering scalings as Eq. (1) interesting. The average solution time, in terms of flips per spin, is shown in Figure 11(a) for $\alpha = 4.1$. Together with Figure 11(b), showing the tendency for the width of the distribution to diminish as $1/\sqrt{N}$, this $1/N$-behaviour implies rather trivial finite size corrections to $P(t_{sol})$. In the Figure we depict the width of the distribution P measured by quantiles instead of the standard deviation, since this is the most sensible measure given the nature of the data.

We also tried to extract the best possible performance of the algorithm as a function of α. Using the data for varying η allows one to extract roughly the optimal values $\eta_{opt}(\alpha)$ which are demonstrated in Figure 12. As can be seen, the data indicate a roughly linear increase of the optimal η with in particular no notice of the approach to α_c or to an algorithm-dependent maximum value α_{max}. The same data can be also utilised to plot, for a fixed N (recall the FMS runs linearly in this regime) the solution time at the optimal noise parameter η. Figure 13 shows that this as expected diverges. Attempts to extract the value

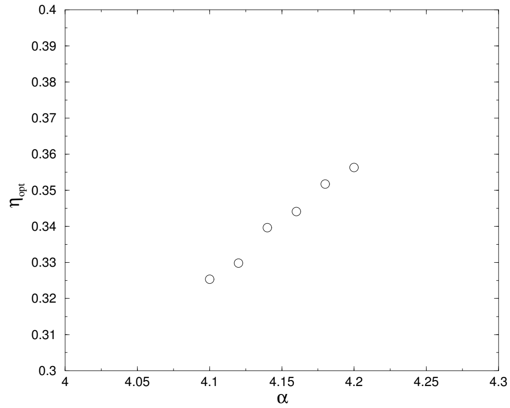

Fig. 12. Optimal η vs. α, for $N = 10^5$.

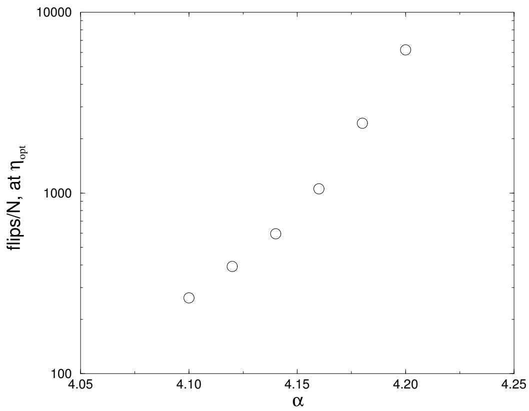

Fig. 13. Solution time at η_{opt} vs. α.

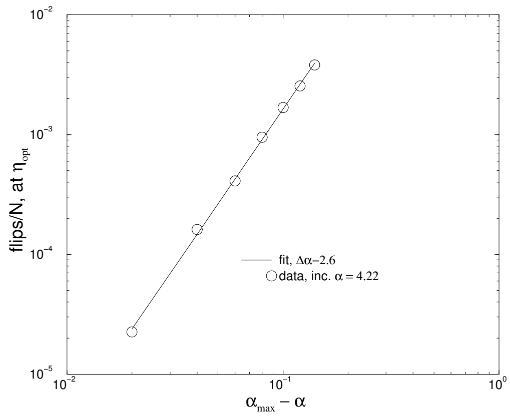

Fig. 14. Solution time t_{sol} at $\eta = \eta_{opt}(\alpha)$ for $N = 10^5$: possible divergence.

α_{max} limiting the linear regime by fitting to various divergences of the kind $t_{sol} \sim (\alpha_{max} - \alpha)^{-b}$ do not allow one to make a statistically reliable conclusion as to whether $\alpha_{max} < \alpha_c$, though . The reason for this is, as far as the data is concerned, the distance of the α studied to any plausible value of α_{max}. See Figure 14.

4 Conclusions

In this paper we have elucidated the behaviour of two focused local search algorithms for the random 3-SAT problem. An expected conclusion is that they can be tuned so as to extend the regime of good, linear-in-N behaviour closer to α_c.

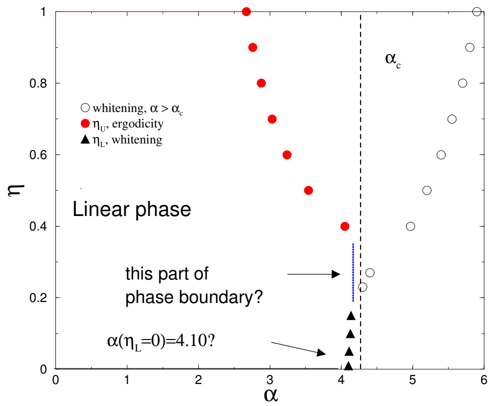

Fig. 15. Proposed phase diagram for FMS. Above a dynamical threshold a metastable state is encountered before a solution is reached. Below a “whitening” threshold a freezing transition happens before a solution or a metastable or stable state is reached. The data points are rough estimates based on a few runs of the FMS algorithm.

An unexpected conclusion is that both algorithms seem to work well also quite close to the critical threshold. Figure 15 proposes a phase diagram. FMS with $\eta = 1$ is just the Random Walk algorithm; hence the first transition point is at $\alpha \approx 2.67$. For larger α there is the possibility of having *two phase boundaries* in terms of the noise parameter. The upper value η_u separates the linear regime from one with too much noise, in which the fluctuations of the algorithm degrade its performance. For $\eta < \eta_l$, greediness leads to dynamical freezing, and though FMS remains ergodic, or able to climb out of local minima, the algorithm no longer scales linearly in the number of variables.

The phase diagram presents us with two main questions: what is the smallest α at which $\eta_l(\alpha)$ starts to deviate from zero? Does the choice of an ideal η_{opt} allow one to push the linear regime arbitrarily close to α_c? Note that the deviation of $\eta_u(\alpha)$ from unity with increasing α could perhaps be analysed with similar techniques as the methods in refs. [3, 25]. The essential idea there is to construct rate equations for densities of variables or clauses while ignoring correlations of neighboring ones. Such analysis, while illuminating, would of course not resolve the above two main questions.

The resolution of these questions will depend on our understanding the performance of FMS or other local search methods in the presence of the clustering of solutions. The range of the algorithm presents us with a fundamental dilemma: though replica methods have revealed the presence of clustering in the solution space starting from $\alpha \approx 3.92$, the FMS works for much higher α's still. The clustered solutions should have an extensive core of frozen variables, and therefore be hard to find. Thus, the ability of FMS to perform well even in this regime, given the right choice of η, tells us that there are features in the solution space and energy landscape that are not yet well understood.

Our numerical experiments also indicate that in the linear scaling regime the solution time probability distributions become sharper ("concentrate") as N increases, implying indeed that the median solution time scales linearly. We cannot establish numerically that this holds also for the average behaviour, but would like to note that our empirical observations from the distributions do not indicate such heavy tails that would contradict this possibility.

Acknowledgements: We are most grateful to Dr. Supriya Krishnamurthy for useful comments and discussions.

References

1. E. Aarts and J. K. Lenstra (Eds.), *Local Search for Combinatorial Optimization.* J. Wiley & Sons, New York NY, 1997.
2. E. Aurell, U. Gordon and S. Kirkpatrick, Comparing beliefs, surveys and random walks. *18th Annual Conference on Neural Information Processing Systems (NIPS 2004).* Technical report cond-mat/0406217, arXiv.org (June 2004).
3. 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.
4. A. Braunstein, M. Mézard and R. Zecchina, Survey propagation: an algorithm for satisfiability. Technical report cs.CC/0212002, arXiv.org (Dec. 2002).
5. A. Braunstein and R. Zecchina, Survey propagation as local equilibrium equations, Technical report cond-mat/0312483, arXiv.org (Dec. 2003).
6. 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.
7. G. Dueck, New optimization heuristics: the great deluge algorithm and the record-to-record travel. *J. Comput. Phys. 104* (1993), 86–92.
8. J. Gu, Efficient local search for very large-scale satisfiability problems. *SIGART Bulletin 3:1* (1992), 8–12.
9. H. H. Hoos, An adaptive noise mechanism for WalkSAT. *Proc. 18th Natl. Conf. on Artificial Intelligence (AAAI-02)*, 655–660. AAAI Press, San Jose Ca, 2002.
10. H. H. Hoos and T. Stützle, Towards a characterisation of the behaviour of stochastic local search algorithms for SAT. *Artificial Intelligence 112* (1999), 213–232.

11. E. Maneva, E. Mossel and M. J. Wainwright, A new look at survey propagation and its generalizations. Technical report cs.CC/0409012, arXiv.org (April 2004).
12. N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller and E. Teller, Equations of state calculations by fast computing machines, *J. Chem. Phys. 21* (1953), 1087–1092.
13. M. Mézard, F. Ricci-Tersenghi, and R. Zecchina, Alternative solutions to diluted p-spin models and XORSAT problems. *J. Stat. Phys. 111* (2003), 505–533.
14. M. Mézard and R. Zecchina, Random K-satisfiability problem: From an analytic solution to an efficient algorithm. *Phys. Rev. E 66* (2002), 056126.
15. 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.
16. 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.
17. 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.
18. G. Parisi, On local equilibrium equations for clustering states. Technical report cs.CC/0212047, arXiv.org (Feb 2002).
19. A. J. Parkes, Distributed local search, phase transitions, and polylog time. *Proc. Workshop on Stochastic Search Algorithms, 17th International Joint Conference on Artificial Intelligence (IJCAI-01).* 2001.
20. A. J. Parkes, Scaling properties of pure random walk on random 3-SAT. *Proc. 8th Intl. Conf. on Principles and Practice of Constraint Programming (CP 2002)*, 708–713. Lecture Notes in Computer Science 2470. Springer-Verlag, Berlin 2002.
21. S. Seitz and P. Orponen, An efficient local search method for random 3-satisfiability. *Proc. IEEE LICS'03 Workshop on Typical Case Complexity and Phase Transitions*. Electronic Notes in Discrete Mathematics 16, Elsevier, Amsterdam 2003.
22. S. Seitz, M. Alava and P. Orponen, Focused local search for random 3-satisfiability. Technical report cond-mat/0501707, arXiv.org (Jan 2005).
23. B. Selman, H. Kautz, 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.
24. B. Selman, H. J. Levesque and D. G. Mitchell, A new method for solving hard satisfiability problems. *Proc. 10th Natl. Conf. on Artificial Intelligence (AAAI-92)*, 440–446. AAAI Press, San Jose CA, 1992.
25. G. Semerjian and R. Monasson, Relaxation and metastability in a local search procedure for the random satisfiability problem. *Phys. Rev. E 67* (2003), 066103.
26. G. Semerjian and R. Monasson, A study of Pure Random Walk on random satisfiability problems with "physical" methods. *Proc. 6th Intl. Conf. on Theory and Applications of Satisfiability Testing (SAT 2003)*, 120–134. Lecture Notes in Computer Science 2919. Springer-Verlag, Berlin 2004.
27. W. Wei and B. Selman, Accelerating random walks. *Proc. 8th Intl. Conf. on Principles and Practice of Constraint Programming (CP 2002)*, 216–232. Lecture Notes in Computer Science 2470. Springer-Verlag, Berlin 2002.