Problem structure and multi-move local search: Criticality and parallelism revisited

Andrea Roli
DEIS
Università degli Studi di Bologna – Italia
andrea.roli@unibo.it
www.lia.deis.unibo.it/∼aro

January 22, 2007
Extended abstract*

1 Introduction

The impact of problem structure on search is a relevant issue in AI research and related areas. This topic has been recently received more attention, due to the following reasons: *(i)* real-world problems are often more difficult to solve than random generated problems of the same size and *(ii)* results obtained by applying statistical mechanics techniques (such as phase transition analysis [3]) have shown a strong correlation between search effectiveness and some critical parameters of the instances at hand.

In this work we investigate the relation between some SAT/MAXSAT instance features and the behavior of local search. We will define structural features on the basis of a constraint graph associated to the instances and in particular we will deeply investigate the impact of the node degree frequency on the behavior of multi-move local search —also known as parallel local search, as more than one local move is applied synchronously at each iteration. This work is inspired by a phenomenon called *criticality and parallelism*, first discovered in combinatorial optimization problems such as the TSP and NK-models [6]. The main result is that increasing parallelism leads to better solutions, but up to a degree at which the solution quality degrades. Moreover, the optimal parallelism is negatively correlated with the system connectivity.

*This paper is a summary of the paper titled *Problem structure and multi-move local search: Results achieved and lines for further research*, available as IRIDIA technical report TR/IRIDIA/2004/24.01 (`www.lia.deis.unibo.it/∼aro/pubs/critpar_sat.ps`) an improved version of which is going to be submitted to a journal.

In this work, we present an empirical analysis to investigate whether and under which hypotheses a similar phenomenon can be observed also in satisfiability problems (both the satisfiability problem —SAT— and the maximal satisfiability problem —MAXSAT). We first address the issue of relating SAT instances with graphs with the aim of defining general structural parameters of the instances. Then we present results of the 'parallelization' of local search on random 3-SAT instances. We show that there exists an optimal value of parallelism which drives local search to achieve an optimal performance in terms of average solution quality. We also observe that the optimal parallelism is negatively correlated with a structural parameter of the graph, namely the average node degree. These results hold both for SAT and MAXSAT instances. Furthermore, we extend our analysis toward structured SAT instances, characterized by an irregular node degree frequency.

2 Criticality and parallelism in combinatorial optimization

The criticality and parallelism phenomenon has been observed in the context of local search algorithms applied to combinatorial optimization problems [6, 5, 4]. Local search strategies start from a candidate solution and iteratively perform small changes to it with the goal to eventually find a (near-)optimal solution. In [6, 5, 4], local search is modified by applying some local moves in parallel.[1]

It has been shown that the effectiveness of these algorithms depends on the parallelism degree τ (number of simultaneous moves): if τ increases, the solution quality also increases up to a maximal point (corresponding to τ_{opt}) at which it starts to decrease. It has also been shown that τ_{opt} is negatively correlated with the average node degree: the higher the degree, the lower τ_{opt}. This result can be explained by observing that the node degree gives a rough estimation of the level of constraint tightness in an instance: The higher the average node degree is, the more tightly the variables are connected. In [6, 5, 4] it is also shown that the optimal parallelism value is associated to a phase transition, hence the term criticality.

3 Structure of Satisfiability Problems

SAT belongs to the class of NP-complete problems [2] and can be stated as follows: given a set of clauses, each of which is the logical disjunction of

[1] The considered algorithms have been implemented sequentially. "Parallel moves" are used with the meaning of "synchronous moves". Anyhow, these results could be beneficial also for implementations on parallel architectures.

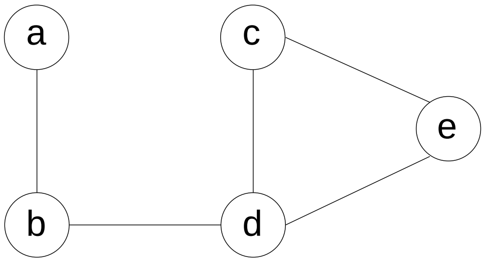

Figure 1: Interaction graph associated to the SAT instance $(a \vee \neg b) \wedge (b \vee d) \wedge (c \vee \neg d \vee \neg e)$. Note that one interaction graph corresponds to a set of formulas, therefore it just represents essential interactions among variables.

$k > 2$ *literals*[2], we ask whether an assignment to the variables exists that satisfies all the clauses. MAXSAT is a NP-hard problem and can be stated as follows: given a set of clauses, the objective is to find an assignment to the variables such that it maximizes the number of satisfied clauses[3].

For SAT instances the *interaction graph* [8] can be defined, in which nodes correspond to variables and an edge connects two nodes if the corresponding variables appear in a same clause (see Fig. 1). Note that other representations are possible, that capture different aspects of problem structure; for instance, a bipartite graph can be defined that represents both literals and clauses as nodes and edges connect literals to the clauses involving them.

The shift from a problem model to its representation via graphs enables us to relate graph properties and characteristics with structural properties of problem instances.

4 Criticality and parallelism in SAT

A phenomenon with analogous characteristics with the one introduced in Sec. 2 has been discovered in parallel local search for SAT [9, 13, 11, 10, 12], in which a parallel version of local search has been applied to both random and structured instances of SAT and its optimization version MAXSAT.

The algorithm used in those studies is a variant of GSAT [14], a greedy local search designed for attacking SAT.[4] GSAT starts from a random assignment and looks for a better one by moving from one state to another

[2] given a boolean variable, the literal is the variable or its negation

[3] If weights are associated to clauses, the objective is to maximize the weighted sum of satisfied clauses. In this case, the problem is called Weighted MAXSAT.

[4] GSAT has been chosen since the previous results on criticality and parallelism were obtained by applying a (quite simple) local search characterized by a strong hill climbing tendency.

one in its neighborhood (defined as the set of states at Hamming distance equal to 1). Given a current state, the next state is chosen by *flipping* the variable that, if flipped, leads to the greatest gradient in the number of satisfied clauses. This algorithm can be easily extended to a multi-move local search by dividing the variables in τ subsets and flipping one variable of each subset according to the GSAT flip criterion. The effect of a variable flip on the objective function value is evaluated as if it were the only one to change, so flips are taken as if they were performed in parallel. In our experiments, the subsets of variables are randomly constructed at the beginning of each local search step. If n is the number of variables, the number of subsets corresponds to the parallelism degree τ and the cardinality of each subset is n/τ[5]. This choice is aimed at having a uniform sample of the average behavior, rather than providing the best subdivision for the instance at hand.

For random instances, it has been experimentally shown that the best performance is achieved with an optimal parallelism degree τ_{opt} that is monotonically non increasing with the average node degree of the graph associated to the instance attacked. A typical example is depicted in Fig.2. The plot shows the average best solution (evaluated as the number of unsatisfied clauses) returned by Parallel GSAT (PGSAT) run with varying number of parallel flips (from 1 to 20 variable flips per move). We can observe that the average best solution curve has a minimum corresponding to a particular value of τ, that is instance dependent. The instance property that affects τ_{opt} is the average node degree of the graph associated to the boolean formula. Fig.3 reports the plot capturing the empirical relation between average node degree and optimal parallelism in the case of random 3-SAT instances. As we can note, τ_{opt} is monotonically non-increasing with the average "connectivity" among nodes.

A very interesting interpretation of the data can be obtained by plotting the optimal parallelism against the average node degree on a log-log plot, as shown in Fig.4. We can note a linear relation[6] that can be interpreted as a sign of a power law. Power law is typical of systems in critical state [1] and it is extremely interesting to note that in [5] it has been conjectured that the optimal parallelism is achieved when the system is in the critical state. Therefore, our results on random 3-SAT can give support to this conjecture.

The analysis of the node degree frequency of the graphs associated with SAT/MAXSAT instances is also particularly informative. The node degree frequency of the random SAT instances is regular (it approximately fits the Gaussian distribution) and the highest peaks are distributed around the mean (see Fig.5). Therefore, the average node degree can be used as a key

[5]More precisely, all subsets have equal cardinality, except for one which contains $\tau + n \ mod \ \tau$ variables.

[6]The correlation coefficient r^2 is 0.8484 and the line slope is -0.85

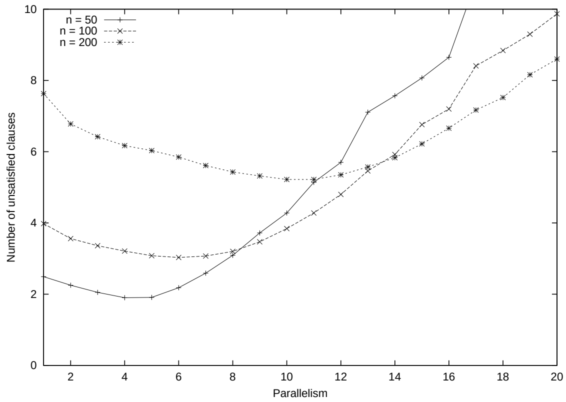

Figure 2: Average error (number of unsatisfied clauses) vs. τ. Random MAXSAT instances with 50,100,200 variables and 218,430,860 clauses respectively. Results are averaged over 10000 runs.

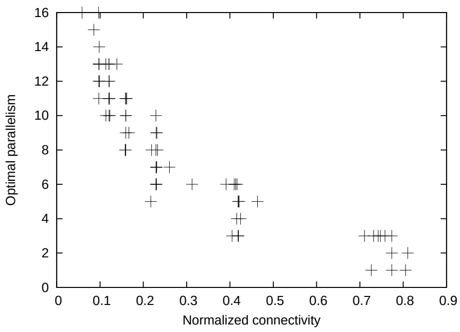

Figure 3: τ_{opt} vs. normalized average node degree.

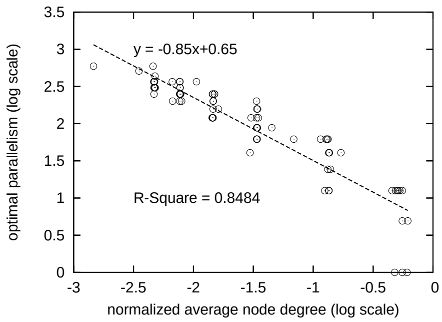

Figure 4: τ_{opt} vs. normalized average node degree in logarithmic scale. Data are fitted by means of a linear regression: $r^2 = 0.8484$ and slope $= -0.85$.

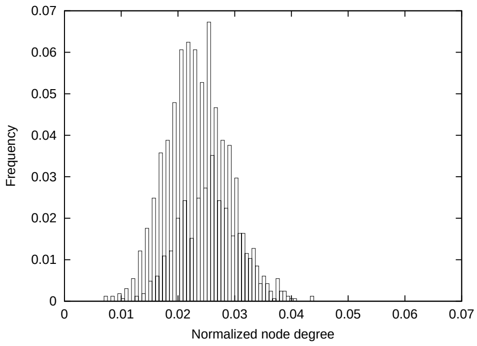

Figure 5: Frequency vs. normalized node degree in a random 3SAT instance with 1650 variables. The normalized average node degree is 0.0239.

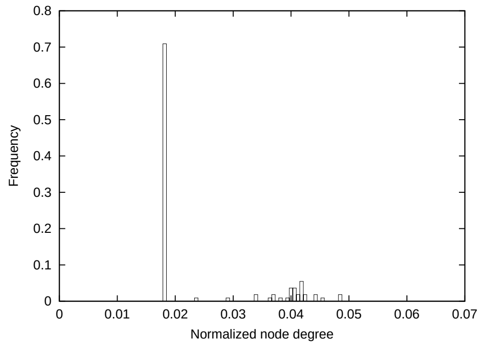

Figure 6: Frequency vs. normalized node degree in the instance *ii16a1* that encodes an inductive inference problem. The instance has 1650 variables and. The normalized average node degree is 0.0239.

parameter for relating instance and τ_{opt}.

Conversely, structured instances, deriving from real-world problems such as circuit testing or planning, have in general a more spread and non-uniform degree frequency. For instance, Fig.6 shows the node degree frequency of a SAT instance encoding an inductive inference problem.[7] This asymmetric and non-Gaussian distribution strongly affects the algorithm performance as a function of τ. In particular, the location of the highest peaks in the frequency turns out to be the most relevant characteristic influencing the optimal parallelism τ_{opt}, i.e., the degree corresponding to the highest peaks is the parameter that mainly affects the actual value of τ_{opt}. For example, we report the case of the behavior of PGSAT on the aforementioned instances, which have the same average normalized node degree $\overline{q} = 0.0239$, but different node degree frequency. Fig.7 and Fig.8 show the average and median error (number of unsatisfied clauses) respectively on the random instance (*3sat1650_q-const*) and the structured one (*ii16a1*) for PGSAT with different values of τ. Results are averaged over 500 trials. We observe that also for the structured instance there exists an optimal value of τ. Nevertheless, despite the fact that the two instances have the same average connectivity, the optimal parallelism is higher for *ii16a1* than for *3sat1650_q-const*. *ii16a1* has, indeed, a high peak close to $0.018 < \overline{q}$.

[7] Instance retrieved from www.satlib.org

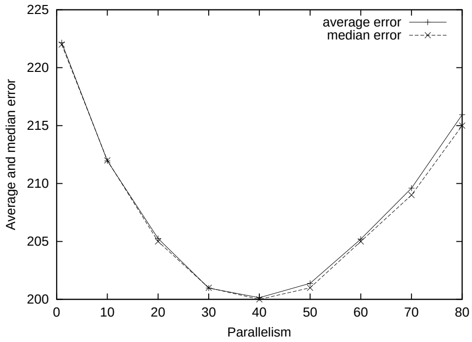

Figure 7: Average and median error against τ for PGSAT on the instance *3sat1650*.

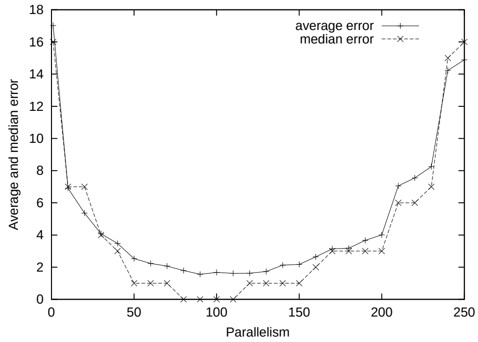

Figure 8: Average and median error against τ for PGSAT on the instance *ii16a1*.

5 Conclusion

In this work, we first showed that for different classes of problems (random and structured SAT, MAXSAT, lattice SAT) there is a value of τ that optimizes the algorithm performance in terms of average solution quality. Furthermore, it has been discovered that, given an interaction graph, its node degree strongly affects the optimal value of τ_{opt}. In case of random instances, the higher the normalized average node degree $\overline{q}$, the lower τ_{opt}. These results are in accordance with the previous results found in the literature. Besides these experiments, the impact of the interaction graph has been investigated by applying PGSAT to structured SAT instances. In our experiments, we have observed that for structured instances there exists an optimal value of τ, as in the case of random and constant degree SAT instances. Moreover, we have observed that τ_{opt} depends on the node degree frequency, and in particular on the frequency peaks.

The relation between node degree and multi-move local search is still to be fully understood. In fact, there are some limitations in the use of the interaction graph and more accurate analyses are required. For example, we believe it is important to consider also the use of different kinds of graphs, such as weighted graphs, to capture the structure of the instances and also a thorough empirical analysis has to be performed on k-SAT instances, with k $\neq$ 3. Furthermore, we can not definitely conclude that the phenomenon is exactly the same found in [7] mainly because in our case it is not possible to define the same parameter defined therein to check for the occurrence of a phase transition.

References

[1] P. Bak. *How Nature Works: The Science of Self-Organized Criticality.* Copernicus, 1996.

[2] M. R. Garey and D. S. Johnson. *Computers and intractability; a guide to the theory of NP-completeness.* W.H. Freeman, 1979.

[3] T. Hogg, B. A. Huberman, and C. P. Williams. Phase transitions and the search problems. *Artificial Intelligence*, 81(1–2), 1996. Special issue on Phase Transitions and Search Problems.

[4] S. A. Kauffman. *At home in the universe.* Oxford University Press, New York, 1995.

[5] S. A. Kauffman and W. Macready. Technological evolution and adaptive organizations. *Complexity*, 26(2):26–43, March 1995.

[6] W. G. Macready and S. A. Kauffman A. G. Siapas. Criticality and parallelism in combinatorial optimization. *Science*, 271:56–59, January 1996.

[7] D. McAllester, B. Selman, and H. Kautz. Evidence for invariants in local search. In *Proceedings of AAAI97/IAAI97*, pages 321–326, Menlo Park, 1997. AAAI Press.

[8] I. Rish and R. Dechter. Resolution versus search: Two strategies for SAT. *J. Automated Reasoning*, 24:225–275, 2000.

[9] A. Roli. Criticality and parallelism in GSAT. *Electronic Notes in Discrete Mathematics*, 9, 2001.

[10] A. Roli. Criticality and parallelism in structured SAT instances. In P. Van Henteryck, editor, *Proceedings of CP02 - Eighth International Conference on Principles and Practice of Constraint Programming*, volume 2470 of *Lecture Notes in Computer Science*, pages 714–719. Springer, 2002.

[11] A. Roli. Impact of structure in parallel local search for SAT. In *SAT 2002 – Fifth International Symposium on the Theory and Applications of Satisfiability Testing*, Cincinnati, Ohio, USA, 2002.

[12] A. Roli. Problem structure and search: Empirical results and open questions. In *Proc. of CP-AI-OR'03 – Fifth Int. Workshop on Integration of AI and OR techniques in Constraint Programming for Combinatorial Optimization Problems*, Montreal (Canada), 2003.

[13] A. Roli and C. Blum. Critical Parallelization of Local Search for MAX–SAT. In F. Esposito, editor, *AI*IA2001: Advances in Artificial Intelligence*, volume 2175 of *Lecture Notes in Artificial Intelligence*, pages 147–158. Springer, 2001.

[14] B. Selman, H. J. Levesque, and D. Mitchell. A new method for solving hard satisfiability problems. In Paul Rosenbloom and Peter Szolovits, editors, *Proceedings of AAAI92*, pages 440–446, Menlo Park, California, 1992. American Association for Artificial Intelligence, AAAI Press.