Penalties may have collateral effects. A MAX-SAT analysis.

Roberto Battiti and **Paolo Campigotto** [1]

Abstract.
Many incomplete approaches for MAX-SAT have been proposed in the last years. The objective of this investigation is not so much horse-racing (beating the competition on selected benchmarks) but understanding the qualitative differences between the various methods. In particular, we focus on *reactive search* schemes where task-dependent and local properties in the configuration space are used for the dynamic on-line tuning of local search parameters and we consider the choice between prohibition-based and penalty-based approaches. To abstract from implementation details we focus on the search trajectory characteristics and study the trade-off between diversification and bias after starting from a local minimizer. We then study the warping effects on the fitness surface of weight-update schemes and the resulting dynamics, through an exhaustive analysis of small MAX-SAT instances, and of the average evolution of individual trajectories. The results are compatible with the conclusion that penalty-based schemes achieve diversification from a starting local optimum through a complex method with global and potentially dangerous collateral effects, while prohibition-based schemes reach comparable or better results in a more direct and controllable manner.

In the final part of this paper, we consider long runs of the complete algorithms on selected MAX-SAT instances, which confirm the competitiveness of prohibition-based reactive approaches.

1 Introduction

Most of the incomplete methods based on stochastic local search (SLS) for MAX-SAT are characterized by a set of parameters whose tuning is crucial for CPU time requirements and solution quality. However, the appropriate tuning depends on both the problem and the current instance being solved, implying costly human intervention. Furthermore, the optimal parameter setting can vary widely in different regions of the configuration space around a given tentative current solution, leading to dynamic adaptive schemes. Reactive search strategies for the on-line dynamic tuning of these free parameters to the current task being solved and to the local characteristics can be used to obtain more robust and efficient techniques [1].

The scope of this paper does not allow a detailed review, see for example [2, 4] for a recent survey of propositional satisfiability and the related constraint programming problem, and [5] for a survey of stochastic local search approaches for MAX-SAT. Let us concentrate on reactive schemes and let us classify them according to the target acted upon during the on-line adaptation. In detail, the reaction can be on the generation of a set of *constraints on the variables* through the *prohibition* of recently-applied moves, or on the *modification of the cost function* guiding the local search. For brevity, the two paradigms will be denoted as *prohibition-based* and *penalty-based*, respectively. The first method, which is an ingredient of what is known as "tabu search," aims at pushing the configuration out from the attraction basin around a local minimizer by temporarily prohibiting some moves which would lead the trajectory back to the starting point. The second method, also termed "dynamic local search," modifies the objective function guiding the search so that a local minimum is raised to encourage the exploration of different areas, see Fig. 1.

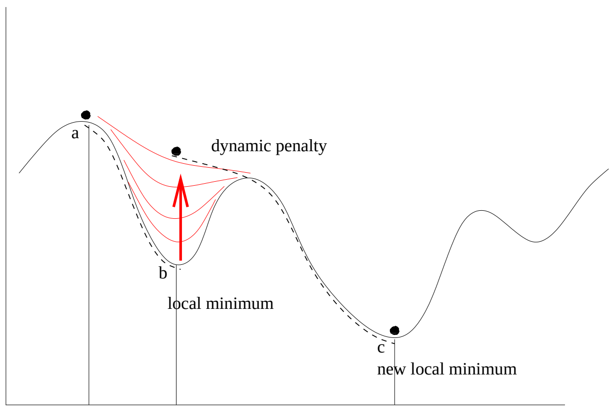

Figure 1. Transformation of the objective function to gently push the solution out of a given local minimum. Note: the intuition can be misleading for dynamically weighted clauses in MAX-SAT.

A second macroscopic difference is given by the selection of the variables considered during each local search step. In the basic schemes (like GSAT), all variables are potential candidates for the next flip, in more recent proposals (like WalkSAT), only the *variables appearing in unsatisfied clauses* are considered for a possible flip (let's call them "unsatisfied variables").

Given the space constraints of this paper we report selected results within an ongoing investigation to compare and identify qualitative differences between search dynamics of prohibition and penalty-based schemes.

[1] Dipartimento di Ingegneria e Scienza dell'Informazione (DISI), University of Trento, Italy, email: {battiti, campigotto}@disi.unitn.it

2 SLS approaches with on-line learning and dynamical systems

Let n and m denote the number of variables and clauses of a given MAX-SAT instance in conjunctive normal form. The simplest cost function guiding SLS algorithms for MAX-SAT is the number of unsatisfied clauses. This function will be denoted as f. Escaping local minima of f in a strategic and intelligent manner can be considered as the underlying motivation of most recent approaches based on stochastic local search. In many cases, local minima are actually large *plateaus* where the basic local search cannot determine the "right" direction to continue and performs a slow random-walk in search of an escape point. More advanced schemes design discrete dynamical systems so that the generated trajectory achieves a more efficient and effective exploration of the fitness surface. In this work we do not consider implementation details (supporting data structures) and CPU times but focus only on the trajectory properties. The fact that appropriate data structures make the advanced schemes fully competitive with the simpler ones has been matter of investigation but it is not included in this work due to space constraints.

A remedy to escape from local minima consists of transforming f into a modified g cost function, therefore *warping* the fitness surface, and generating a new direction of movement. This cost function modification may look at the *internal structure* of the current solution, not simply at the number of satisfied clauses. In [1] one exploits *non-oblivious* cost functions, which measure the *degree* of satisfaction of each clause by counting the number of matched literals. Aiming at a redundant satisfaction eliminates the difficulty in selecting among seemingly similar situations and may eventually permit to flip a variable to satisfy a new clause, without losing any already satisfied clause.

Another approach to escape from local minima or plateaus of f is given by *dynamic local search*, that relies on a dynamically weighted version of the oblivious function. The works in [8, 10] use dynamic weights to encourage the satisfaction of "more difficult" clauses. Clause weighting is motivated in [8] as a "breakout method for escaping from local minima", in [10] as a way to "fill-in" local minima. See for example [18, 15] and the contained references for some recent work in this area.

A starting point of this work is [17] where the authors investigate the dynamic warping of the search space caused by dynamic weight penalties, fail to find evidence that warped landscapes represent accumulated knowledge about the search space [3] and clarify that the "hole-filling" analogy can be deceiving by presenting a toy example showing the global and potentially detrimental side-effects, also hinted in [8]. Their empirical investigation shows that warping algorithms mainly serve as a diversification mechanism, which allows the search process to effectively overcome stagnation due to local minima and plateaus.

The purpose of this work is to continue the investigation through additional means:

Exhaustive analysis of warped landscapes Small (but non-trivial) instances of MAX-SAT are subjected to an exhaustive analysis of the modification (warping) effects by examining the local minima canceled, produced, and maintained after updating weights.

Diversification-Bias analysis It is suggested in [1] that Pareto-optimal points on D-B plots of basic versions of SLS schemes have an empirical predictive power for the overall success of the methods. This hypothesis is investigated for skeletal versions of prohibition- and penalty-based schemes.

Sample trajectories analysis While D-B plots summarize snapshots of the search after short periods starting from a local minimum event, and the exhaustive analysis describes the overall modification of the fitness surface, the analysis of sample trajectories produces additional information about the dynamical evolution of the search.

3 Exhaustive analysis of warped landscapes

First we perform an experimental analysis on a single unsatisfiable MAX-SAT instance formed by 20 variables and 110 clauses, close to the "satisfiability threshold region" [13].

We consider a prototypical and simplified version of the weight-update approach: as soon as the first local minimum (FLM) point for the "standard" f function is encountered, the weights of the currently unsatisfied clauses are increased by a fixed quantity Δw. An exhaustive analysis of the search space is then performed, showing the difference between the original and the warped fitness landscape generated by the weight update. The number of local minima of the landscape at different Hamming distances from the FLM point are counted. This classification allows to understand if the effects of the weight update are local or global, i.e., if the changes of the fitness landscape concentrate in the neighborhood region surrounding the FLM point or affect the whole search space. Due to the weight update operation performed when visiting the FLM point, new local minima can be generated in the search space (Fig. 2), while, at the same time, "old" local minima (i.e., local minima in the original landscape) may be canceled (Fig. 3). By definition, a local minimum disappears as soon as at least one improving move is available in its neighborhood.

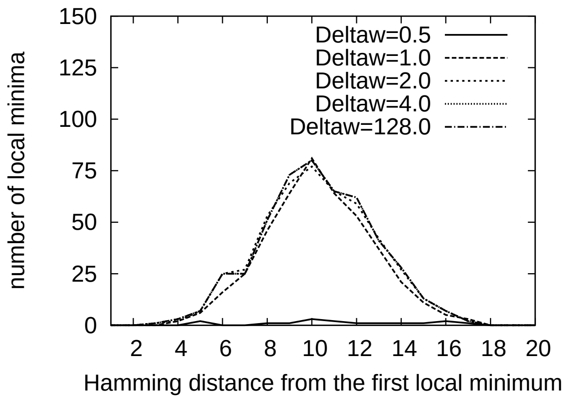

Figure 2. Distribution of the newly generated local minima over different warped landscapes. The curves describe the landscape generated by a weight increase operation with different Δw values ($0.5, 1, 2, 4, 128$). The curve for $\Delta w = 0.5$ is at the bottom of the Fig.

The absolute numbers of local minima generated or canceled is bigger in the region around Hamming distance 10 from the FLM point. One suspects that the numbers are related to the original number of local minima present at specific distances, which is in turn related to the total number of binary strings at specific distances, a distribution peaked at distance $n/2$ with simple counting arguments.

Furthermore, the observed effects over the search landscape are not directly related to Δw: as soon as the Δw value is bigger than 1.0, very similar curves are obtained. The warped landscapes generated by the considered weight update values ranging from 4.0 to

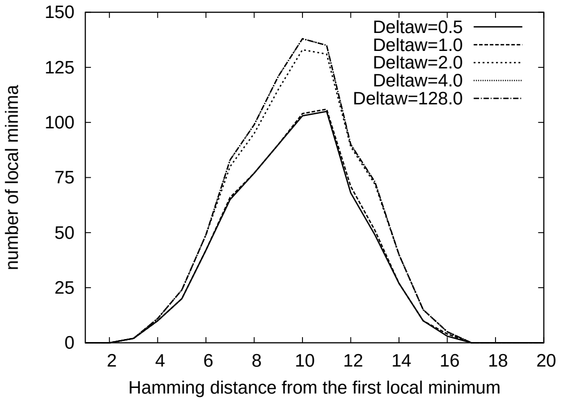

Figure 3. Distribution of the deleted local minima over different warped landscapes. The curves describe the landscape generated by a weight increase operation with different Δw values $(0.5, 1, 2, 4, 128)$. The curve for $\Delta w = 4.0$ overlaps with the curve for $\Delta w = 128.0$, while the curve for $\Delta w = 0.5$ and $\Delta w = 0.5$ are very close.

128.0 delete and generate exactly the same number of local minima at specific Hamming distances from the FLM point: the weighted clauses become so important that the effect of the other clauses is negligible.

Finally, the number of the canceled local minima is bigger than the number of the newly generated local minima.

Fig. 4 considers ratios instead of absolute numbers. Ratios underline that the deletion process of "old" local minima acts in a rather uniform way over the whole search landscape: it does not depend on the Hamming distance from the FLM point. Therefore, the changes affecting the search landscape in no way can be considered a localized effect.

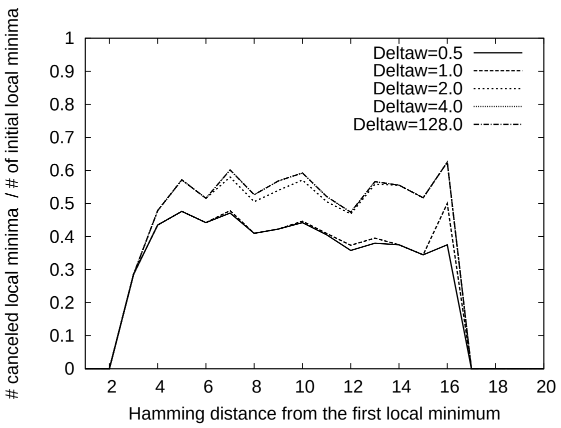

Figure 4. Fraction of canceled local minima over the search landscape. The curves are for the different Δw values $(0.5, 1, 2, 4, 128)$

The results clearly confirm the hints and the anecdotic evidence of previous researchers: weight updates caused by a specific local minimum encountered along the trajectory have a huge global effect: to escape from a single local attractor one is modifying in a radical manner also the value of configurations which are very far from the local minimum. In addition, a very large fraction of the initial local minima are canceled, about 50% in the given example, after an update caused by a single local minimum.

To measure the quality of the warped fitness landscape w.r.t. the original landscape, the quality of the deleted/generated local minima is also taken into account. The quality of the local minima is measured in terms of the cost function f, which counts the number of unsatisfied clauses. The smaller is the f value, the better is the quality of the local minimum, as it can be considered a better "approximation" of the global minima.

In particular, do weight-update schemes generate a "better" landscape, i.e., a landscape allowing better performance for local search strategies? This would be the case if poor quality local minima are ironed out but good quality ones are preserved. Furthermore, is the deletion/generation of the local minima related to their quality?

The weight-update mechanism aims at raising up (i.e., deleting) the local minimum which the algorithm is trapped in, allowing to escape from it (see Fig. 1). Let's call this local minimum the *current* local minimum. (Note that in our experiments the *current* local minimum is the FLM point).

There is a popular belief that the weight-update mechanism tends to delete low quality local minima in addition to the current one, obtaining a "cleaner" fitness surface. As a results, the warped surface should speed up the search of the global minima.

However, for the instance considered in the exhaustive analysis, we show this is not the case. Fig. 5 and 6 show the mean quality of the generated and deleted local minima w.r.t. the mean quality of the local minima in the original fitness surface. The different plots are for different weight update values shown in the labels. The plots for the weight update values 8.0, 16.0, 32.0 and 128.0 are equal to the plot for $\Delta w = 4.0$.

The quality of the newly generated local minima is worse than the quality of the deleted ones. Furthermore, for all the weight update values considered, the mean quality of the local minima in the warped surface never improves w.r.t. the mean quality of the local minima in the original surface (Table 1). In most of the cases, it is worse.

Δw	LM number	mean LM quality
0.0	1629	6.7114
0.01	958	6.5417
0.1	957	6.5423
0.5	974	6.5533
1.0	1383	6.8886
2.0	1244	7.0506
4.0	1227	7.1035

Table 1. Snapshot of the fitness surface after the weight-update operation. The first column indicates the Δw value, while the second and third ones show the number of local minima and their mean quality, measured in terms of the cost function f. The first line (weight update equal to 0.0) indicates the original "fitness surface".

Again, the strong global effect is observed: the quality of deleted/generated local minima does not depend on the Hamming distance from the FLM point.

In this experiment the *global* minima of the studied MAX-SAT instance (corresponding to 3 points with score function equal to 2) are not affected by the collateral effect of the weight update, i.e., they are still *global* minima in the warped surface.

The analysis performed shows that the warped surface obtained

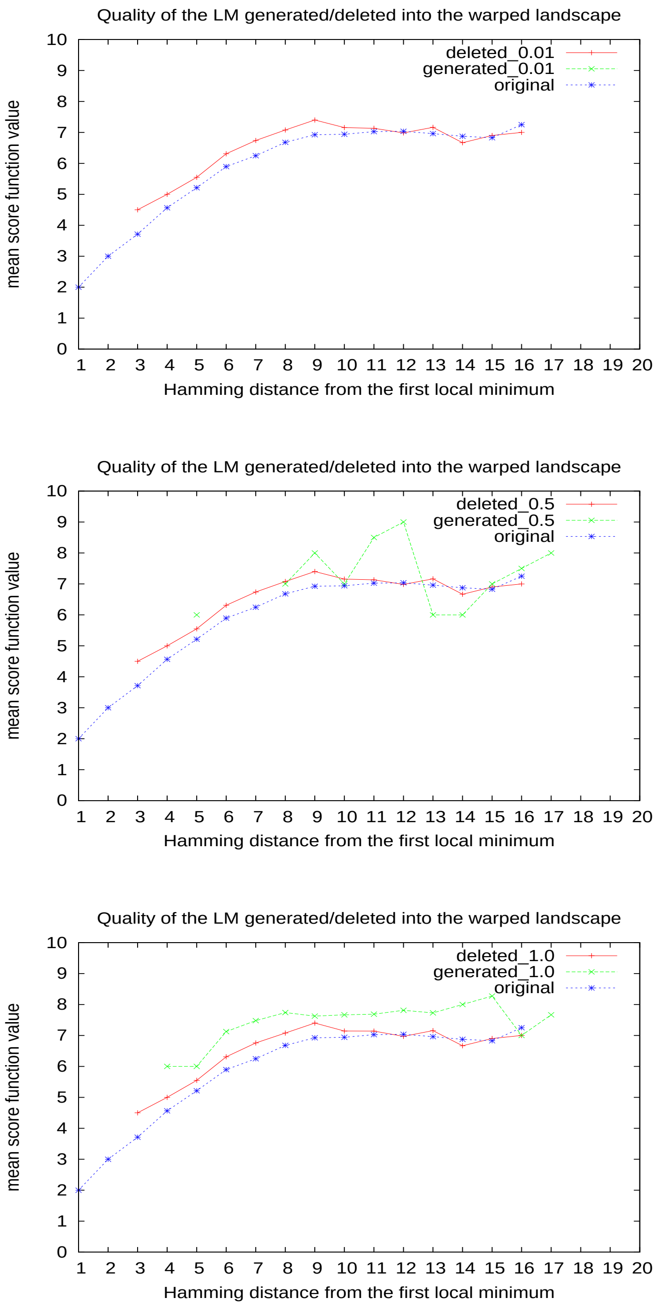

Figure 5. The mean quality of the generated/deleted local minima for the weight update values 0.01, 0.5, 1.0. The curve labeled as "original" shows the mean quality of the local minima in the initial fitness surface.

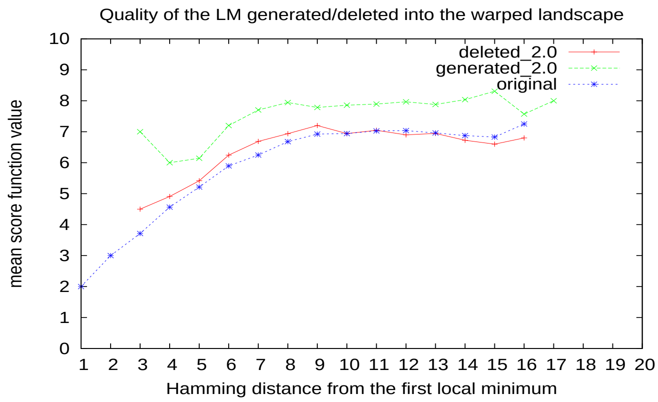

(a)

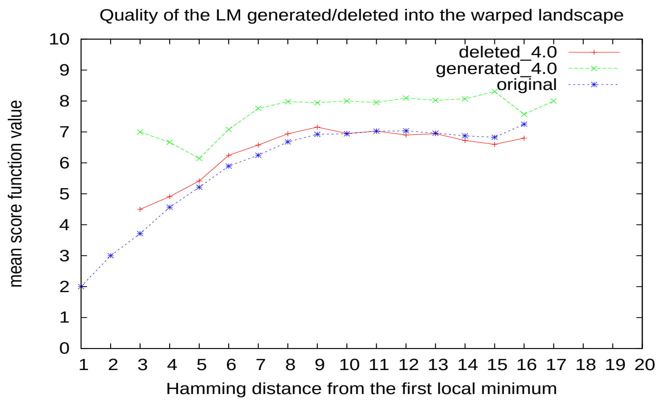

(b)

Figure 6. The mean quality of the generated/deleted local minima for the weight update values 2.0, 4.0. The curve labeled as "original" shows the mean quality of the local minima in the initial fitness surface.

in all the experiments does not correspond to one of a better quality for the local search techniques. The long runs of complete penalty-based algorithms reported in the last part of the paper fully validate this observation.

To verify that our observations are not biased by the MAX-SAT instances benchmark selected, we now consider a small hand crafted MAX-SAT instance obtained from the SAT 2005 competition, with 24 variables and 61 clauses, and repeat the exhaustive analysis of the search space. Fig. 7, 8 and 9 show the distribution of the deleted/generated local minima and the ratio among the deleted and the initial local minima, respectively.

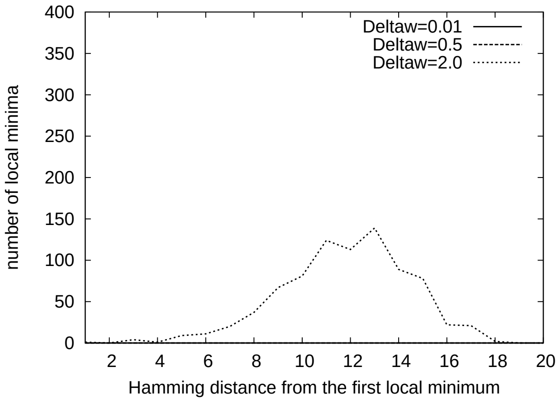

Figure 7. Distribution of the newly generated local minima over different warped landscapes. The curves describe the landscape generated by a weight increase operation with different Δw values $(0.01, 0.5, 2.0)$. For $\Delta w = 0.01$ and $\Delta w = 0.5$ no new local minima are generated.

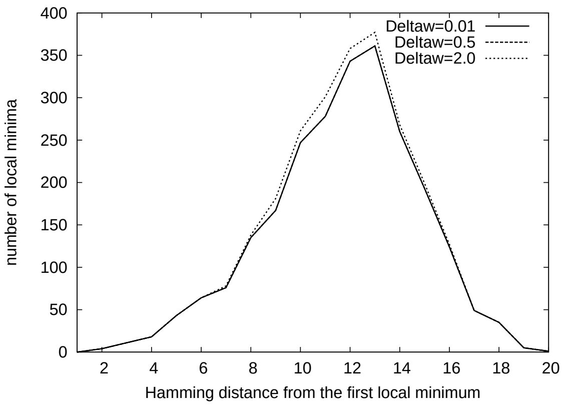

Figure 8. Distribution of the canceled local minima over different warped landscapes. The curves describe the landscape generated by a weight increase operation with different Δw values $(0.01, 0.5, 2.0)$

Fig. 10 compares the mean quality of the generated/deleted local minima w.r.t. the quality of the initial local minima for weight update values equal to 0.01, 0.5, and 2.0. The warped surface obtained for the weight update values 4.0, 8.0, 16.0, 32.0, 64.0 and 128.0 is equal to the case $\Delta w = 2.0$.

Table 2 summarizes the changes among the original and the warped fitness surface for the various experiments.

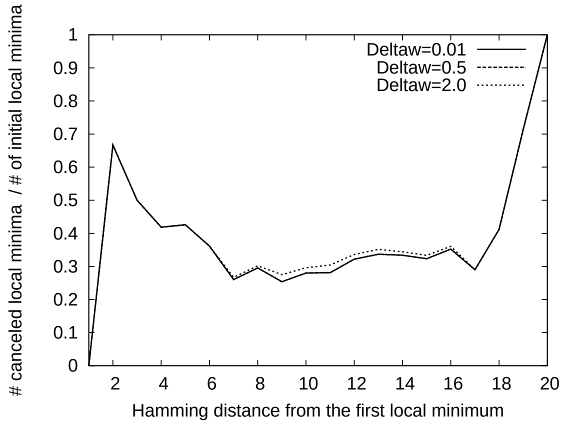

Figure 9. Fraction of canceled local minima over the search landscape. The curves are for the different Δw values $(0.01, 0.5, 2.0)$

Δw	LM number	mean LM quality
0.0	7756	3.7748
0.01	5341	3.5967
0.1	5341	3.5967
0.5	5341	3.5967
1.0	6139	3.7730
2.0	6056	3.7836

Table 2. Snapshot of the fitness surface after the weight-update operation. The first column show the weight update considered, while the second and third ones show the number of local minima and their mean quality, measured in terms of the cost function f. The first line (weight update equal to 0.0) indicates the original "fitness surface".

The results on the crafted instance confirm the observations for the random instance:

- the distribution and the quality of the deleted/generated local minima do not depend strongly on the Hamming distance from the FLM point;
- the number of the deleted local minima is bigger than the number of newly generated local minima;
- the mean quality of the local minima in the new fitness surface is never strongly improved (in some cases, it is even worse).

Again, the penalty-based approach exhibits a global and potentially undesiderable side effect.

Furthermore, in this case the *global* minima are affected by the weight- update operation. In the original fitness surface there are 134 *global* minima; for the weight update values bigger than 2.0 considered, 12 of them are deleted: over the warped surface they are not even a local minimum. For the weight update values smaller than 1.0, one initial global minimum becomes a local minimum point in the warped surface, while 11 initial global minima do not remain local minima.

4 Diversification-bias analysis

We follow the diversification-bias empirical analysis ("D-B plots") proposed in [1] where the authors conjecture that the dominant

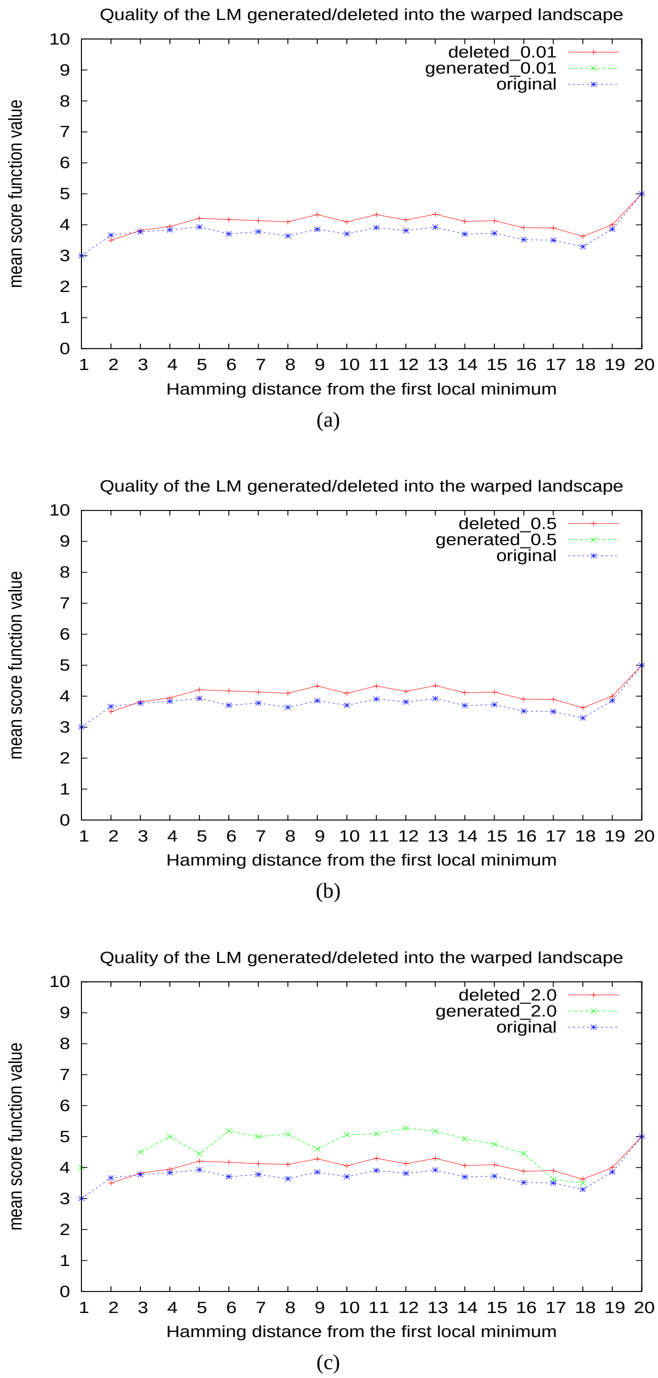

Figure 10. The mean quality of the generated/deleted local minima for the weight update values 0.01, 0.5, 2.0. The curve labeled as "original" shows the mean quality of the local minima in the initial fitness surface.

Pareto-optimal points on D-B plots of basic versions of SLS schemes have an empirical predictive power for the overall success of the methods. The metric used to measure the quality of the visited points (or, simply, the *bias* of the algorithm) is their average cost function value, while the diversification is measured via the average Hamming distance reached in short runs starting from a local optimum. A scheme (characterized by the choice of algorithm and parameter, Δw or prohibition T) is Pareto-optimal if there is no other scheme reaching both a higher diversification *and* a better bias. The prohibition scheme (GSAT/tabu) acts according to the very simple rule: after a single bit (truth value) is changed, it cannot be changed again for the next T iterations. Among the admissible bits (the ones which can be changed), one leading to the best Δf value is chosen randomly among the possible ties.

By using the D-B plots, we want to understand how the warped landscapes generated by the weight-update schemes affect the performance of the dynamic local search (DLS) algorithms, considering both bias and diversification. In particular, we compare the results for the weight-update method with the results of the prohibition-based approach.

All runs of the algorithms considered proceed as follows: as soon as the first local optimum for the "standard" f function is encountered, it is stored and the algorithm is then run for additional $4 * n$ iterations. The final D-B values averaged over 500 tests are reported. For each test we identify the first local minimum via the GSAT algorithm, and then, depending on the different test, we run one among GSAT/tabu and the weighted version of GSAT, starting from the the discovered local minimum.

The tests presented in this work are dedicated to selected MAX-3-SAT instances defined in [7]. In detail, if $n : m$ identify variables and clauses, 50 instances for the 20:110 cases have been randomly generated. The different algorithms are run for the different instances, for a total of 500 tests. The average results are presented.

First, we evaluate the GSAT/tabu method based on fixed prohibitions. Then we study the performance of a simplified version of the weighted GSAT algorithm. Initially all clause weights are equal to one and, once the first local optimum is encountered, the weights of the currently unsatisfied clauses are increased by a fixed quantity Δw. Let's note that the increment of the weights is performed only once: for the subsequent local minima, weights remain fixed.

Fig. 11 shows the results obtained by running the considered algorithm for $4 * n$ steps after the first local minimum discovered by the GSAT algorithm for the same SAT instances. The labels for the curve named *gsatWeighted* represent the different Δw values considered. The value 0 for the "gsatWeighted" curve represents the case of the original GSAT algorithm [12]. The curve named *gsatTabu* is labeled with the values for the fractional prohibition T_f, given by the prohibition parameter divided by n.

The bias is plotted as *difference* w.r.t. the starting f value at the local minimum (good values are therefore at the bottom), the Hamming distance is divided by the number of variables n. One notes that small prohibitions values lead to bias levels comparable with the penalty-based scheme, but they allow for a bigger diversification. This result is confirmed by the repetition of the experiment over the 500:5000 instances, which shows a bigger set of Pareto-optimal points for the prohibition-based scheme (Fig. 12).

In Sec. 6 we consider the original complete schemes for penalty-based and prohibition-based approaches, including the reactive and dynamic versions and analyze the average f values obtained for long runs. The D-B plots Pareto-optimal points are indeed accurate predictors of performance over the long runs.

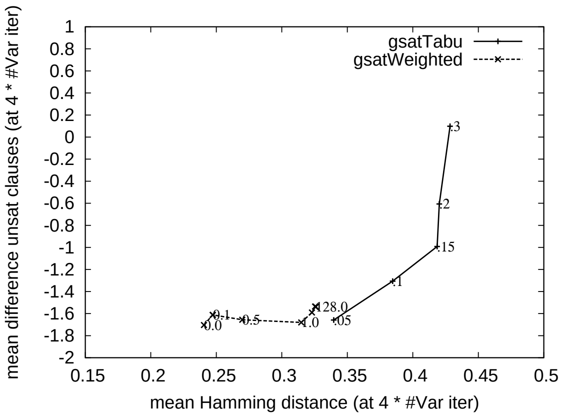

Figure 11. Diversification-bias plots of prohibition-based and penalty-based strategies (20:110 MAX-SAT instance). The curve named *gsatTabu* is labeled with the values for the fractional prohibition T_f, the curve named *gsatWeighted* is labeled with the Δw values for the weights update (points for $\Delta w = 8, 16, 32, 64, 128$ are at a very similar position).

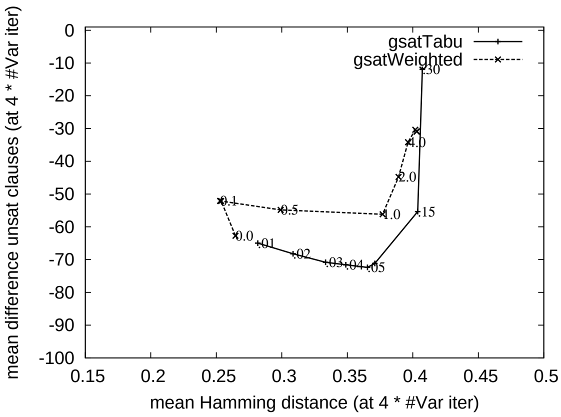

Figure 12. Diversification-bias plots of prohibition-based and penalty-based strategies (500:5000 MAX-SAT instance).

5 Sample trajectories analysis

We now consider the impact of the penalty-based and of the prohibition-based mechanisms on the dynamical evolution of the search trajectories. In particular, do the changes on the landscape caused by the weight-update strategy significantly affect the behavior of a dynamic local search algorithm? If yes, is the biased behavior observed while visiting the region surrounding the first FLM point or over the whole search landscape? Furthermore, we ask whether the prohibition-based strategy, that does not modify the fitness surface, performs a similar function but with a more direct manner. Let's consider again the simplified prohibition-based mechanism and the weight-update strategy described so far. We execute a single short run of both methods starting from a FLM point on the 20:110 instance considered (the first point encountered by local search starting from an initial random configuration), and measure the diversification level reached in both cases (Fig. 13 and 14). We repeat the experiment using different values of the fractional prohibition T_f and of the Δw quantity. The trajectories of the prohibition-based scheme

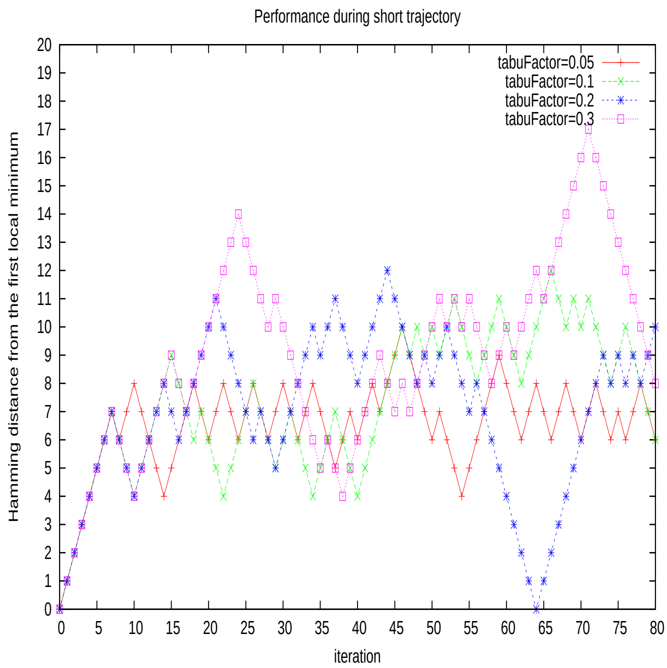

Figure 13. Evolution of the Hamming distance of single short trajectories of the prohibition-based strategy .

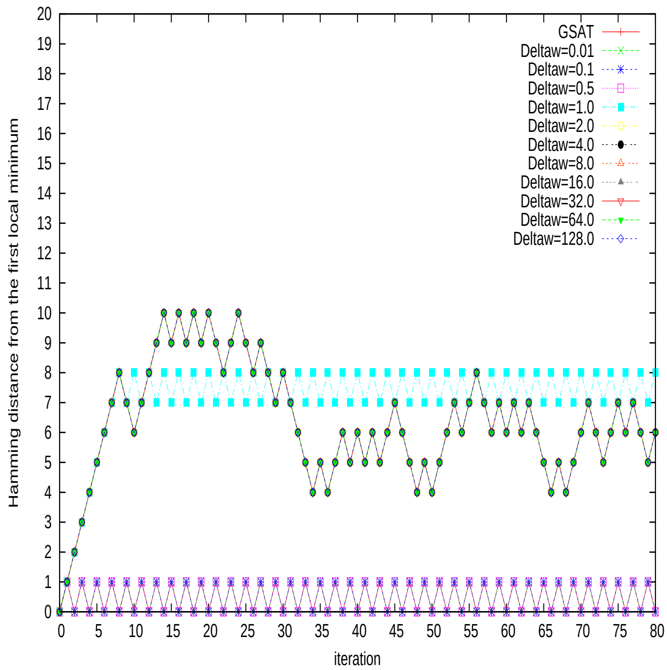

Figure 14. Evolution of the Hamming distance of single short trajectories of the penalty-based strategy. The curves for $GSAT$ and Δw values smaller than 1.0 overlap (bottom of the Fig.), and the curves for Δw values bigger than 1.0 are the top ones overlapping.

show an initial Hamming distance growing linearly up to distance $T+1$, a deterministic effect cause by the prohibition mechanism, followed by a diverse and randomized exploration of the search space. The average deviation of the distances tends to grow with the larger prohibition values. No evidence of entrapment is shown in the later steps.

The situation is qualitatively different for the penalty-based scheme. For small Δw values (values smaller than 1.0, in this case), the trajectory shows a cycle of length 2. For bigger values, the algorithm escapes from the attraction basin of the FLM point, but it is eventually stuck at another local minimum. For the 1.0 value, the algorithm cannot escape from this second local minimum. Furthermore, for all the values bigger than 1.0 the same trajectory is observed. This is not surprising, as in the exhaustive analysis we observed that the number of canceled/generated local minima is the same in all cases. To be fair, let's note that the observed cycles may be avoided by considering more complex weight-update based mechanism, that perform the weight update for each local minimum encountered during the trajectory.

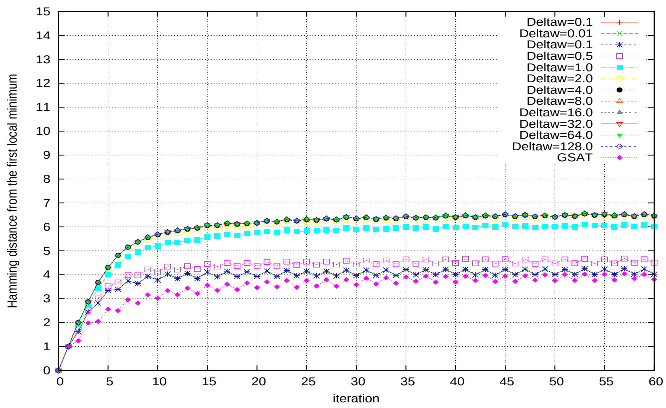

Figure 16. Performance of penalty-based strategy during short trajectories. The values are averaged over 1000 trajectories starting at a different initial FLM point.

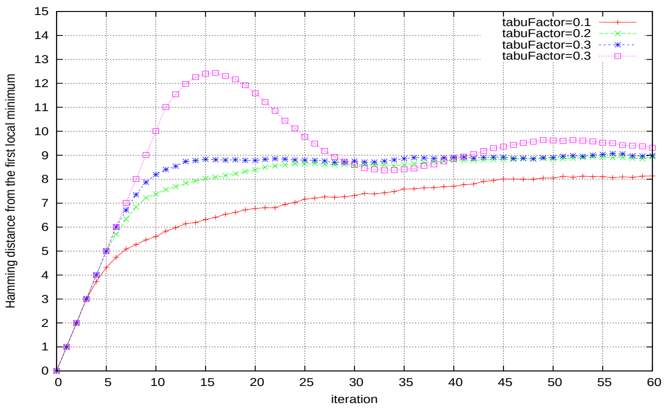

Figure 15. Performance of prohibition-based strategy during short trajectories. The values are averaged over 1000 trajectories starting at a different initial FLM point.

To validate our initial observation of sample trajectories, we now consider the performance of penalty-based and prohibition-based approaches averaged over 1000 runs performed starting from 1000 different initial local minima. Fig. 15 show that during the first iterations of the prohibition-based approach the diversification strictly increases, as expected, and that, eventually, it tends to converge to a common value. The memory about the initial local minimum is effectively lost and exploration proceeds without hindrance. Furthermore, larger values for the prohibition T lead to a bigger initial diversification. Fig. 16 clearly indicates that the performance in terms of diversification of the penalty-based approach is worse than that of the prohibition-based strategy: smaller Hamming distances are reached and the effect shows a fragile dependence on the Δw values, which can be compared to the rather similar behavior of different T values after runs of comparable length (60 iterations in our case). As suggested by intuition, a better diversification is reached with the bigger Δw values. In particular, the worst performance is reached by the GSAT algorithm, that operates over the non-weighted f function.

6 Experiments on long runs

Even if this work does not target the horse-racing point of view, to validate the results of the exhaustive analysis and of the D-B plots experiments, we show the MAX-SAT results reached by the penalty-based and the prohibition-based approaches. In particular, we consider the 500:5000 and the 300:1500 MAX-SAT benchmarks and execute the following SLS approaches:

- GSAT [12], a basic local search greedy strategy guided by the score function f, that simply counts the number of unsatisfied clauses;
- GSAT/tabu [14], which enriches the GSAT algorithm via a prohibition-based search criterion;
- WalkSAT/SKC [11], the ancestor of the WalkSat family. It randomly alternates between greedy minimizing moves and random noisy moves. The moves of both kinds act on the variables appearing in unsatisfied clauses;
- WalkSAT/tabu [6], that adopts the same score function and the same variables selection mechanism of the WalkSAT/SKC algorithm, complemented by tabu search;
- H-RTS, a Hamming-based reactive tabu search algorithm, that dynamically adapts the prohibition parameter during the search;
- AdaptNovelty$^+$ [16], that exploits the concept of variable "age" and uses the same scoring function of GSAT. The variable age can be considered a sort of soft prohibition of recently-changed variables in the case of ties. The prefix *"Adapt"* underlines a reactive behavior, that dynamically adjusts its internal parameters;
- Scaling and Probabilistic Smoothing (SAPS) [18], an accelerated version of the Exponentiated Subgradient algorithm [9] based on dynamic penalties, and a reactive version thereof called RSAPS.

For brevity we report here only the average results (10 runs with different random seeds for each of the 50 instances) as a function of the number of iterations (flips). The user of SLS algorithms is typically interested in the number of iterations required by each algorithm to reach the desired results, or, at least, a good quality approximation. As predicted by the previous diversification-bias analysis and according to the exhaustive search experiments performed, the curves in Fig. 17 confirm a clear superiority of the prohibition-based techniques with respect to the penalty-based approaches. The error bars are not shown on the plots to avoid cluttering. Among all the possible values for the tabu parameter of the WalkSAT/tabu algorithm, we plot the case where the fractional prohibition T_f is 0.01, as with this setting we obtain the best performance over the considered benchmark. The same for the GSAT/tabu algorithm, whose curve is drawn for the optimal T_f value 0.05 over our benchmark set.

The SAPS parameters have been set to the default values, without attempting any extensive optimization. Preliminary tests obtained changing the values did not lead to significant improvements.

With this optimal setting, the GSAT/tabu algorithm reaches eventually a performance equivalent to that of H-RTS, even if its performance is inferior in the initial phase. This result clearly indicates that parameters setting is crucial for the algorithms performance: not only H-RTS reaches results comparable to the ones with a fixed and optimal T_f, but it actually improves on these because of the dynamic on-line adaptation. This observation is emphasized by the curves for SAPS and RSAPS. They confirm the effectiveness of the reactive approach, that obtains better results while, at same time, allowing to avoid the manual tuning. The SAPS parameters have been set to the default values, without attempting any extensive optimization. Preliminary tests obtained by changing the values did not lead to significant improvements.

Finally, the curve for H-RTS shows the effectiveness of the NOB search to rapidly discover good local optima.

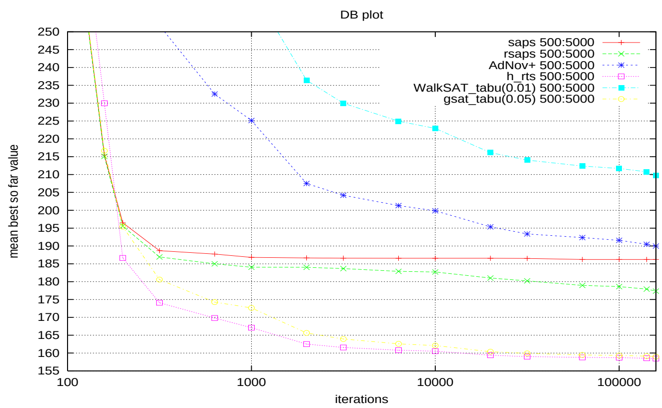

Figure 17. The mean best so far bias value reached by the SAPS, RSAPS, AdaptNovelty$^+$ and H-RTS algorithms.

Fig. 18 shows the behavior of the same algorithms in a scenario closer to the satisfiability threshold (the clauses/variables ratio of the 300:1500 tasks is 5) and to the small 20:110 MAX-SAT instance considered for the exhaustive analysis. The results of the previous long runs experiment are confirmed, apart from the competitive performance of AdaptNovelty$^+$ which eventually duplicates H-RTS' performance although with a much slower start.

Let us note that many of the considered techniques have been proposed for SAT and one may argue that a direct comparison with MAX-SAT algorithms such as H-RTS is not fair. On the other hand, the underlying logic of the methods is always based on maximizing the number of satisfied clauses, which is an argument in favor a direct comparison, in particular for adaptive techniques. In any case, this issue will be explored further in the future.

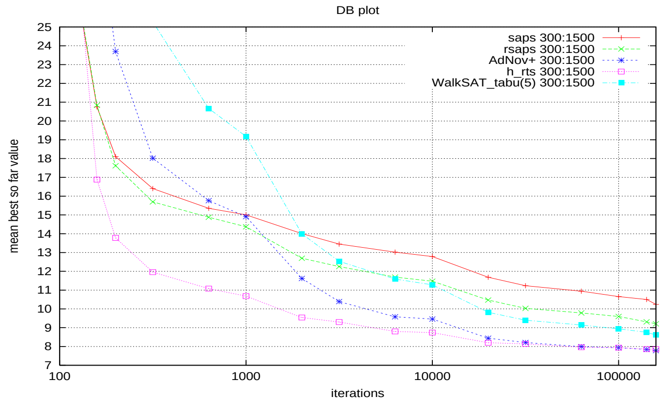

Figure 18. The mean best so far bias value reached by the SAPS, RSAPS, AdaptNovelty$^+$ and H-RTS algorithms on 300:1500 instances.

7 Conclusion

We presented some selected results of an ongoing comprehensive evaluation of alternatives design strategies for MAX-SAT algorithms base on stochastic local search. In particular, we focused on studying the qualitative differences of the dynamics caused by prohibition- and penalty-based schemes, by exhaustively analyzing the warped landscape of small problems, by measuring the diversification and the bias after starting from local minima, and the average behavior of sample trajectories.

The results confirm the hypothesis that penalty-based modifications of the search landscape have global side-effects with a potentially disturbing influence on the search trajectory. The real advantage of these schemes appears to be the fact of forcing the trajectory to abandon an area around a local optimizer to avoid confinement.

On the other hand, a very similar "escaping" effect can be obtained by direct prohibition-based schemes, which do not require the addition of explicit forgetting schemes as a cure to the potentially harming side-effects and are more amenable to explanation. We are aware of the preliminary and in part controversial nature of this investigation, which motivates additional work to further substantiate this hypothesis, when both dynamical system properties and the final competitiveness of the implemented schemes are considered.

REFERENCES

[1] R. Battiti and M. Protasi, 'Reactive search, a history-sensitive heuristic for MAX-SAT', *ACM Journal of Experimental Algorithmics*, **2**(ARTICLE 2), (1997). http://www.jea.acm.org/.

[2] L. Bordeaux, Y. Hamadi, and L. Zhang, 'Propositional Satisfiability and Constraint Programming: A comparative survey', *ACM Computing Surveys (CSUR)*, **38**(4), (2006).

[3] J. Frank, 'Learning short-term weights for GSAT', in *Proc. of the Fifteen INTERNATIONAL JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE*, volume 15, pp. 384–391. LAWRENCE ERLBAUM ASSOCIATES LTD, USA, (1997).

[4] C.P. Gomes, H. Kautz, A. Sabharwal, and B. Selman, *Handbook of Knowledge Representation.*, chapter Satisfiability solvers, Elsevier, 2008, in press.

[5] H. H. Hoos and T. Stuetzle, *Stochastic Local Search: Foundations and Applications*, Morgan Kaufmann, 2005.

[6] D. McAllester, B. Selman, and H. Kautz, 'Evidence for invariants in local search', in *Proceedings of the national conference on artificial intelligence*, number 14, pp. 321–326. John Wiley & sons LTD, USA, (1997).

[7] D. Mitchell, B. Selman, and H. Levesque, 'Hard and easy distributions of SAT problems', in *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*, pp. 459–465, San Jose, Ca, (July 1992).

[8] P. Morris, 'The breakout method for escaping from local minima', in *Proceedings of the national conference on artificial intelligence*, number 11, p. 40. John Wiley & sons LTD, USA, (1993).

[9] D. Schuurmans, F. Southey, and R.C. Holte, 'The exponentiated subgradient algorithm for heuristic boolean programming', in *Proceedings of the international joint conference on artificial intelligence*, volume 17, pp. 334–341. Lawrence Erlbaum associates LTD, USA, (2001).

[10] B. Selman and H.A. Kautz, 'An empirical study of greedy local search for satisfiability testing', in *Proceedings of the eleventh national Conference on Artificial Intelligence (AAAI-93)*, Washington, D. C., (1993).

[11] B. Selman, H.A. Kautz, and B. Cohen, 'Noise strategies for improving local search', in *Proceedings of the national conference on artificial intelligence*, volume 12. John Wiley & sons LTD, USA, (1994).

[12] B. Selman, H. Levesque, and D. Mitchell, 'A new method for solving hard satisfiability problems', in *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*, pp. 440–446, San Jose, Ca, (July 1992).

[13] Bart Selman, David G. Mitchell, and Hector J. Levesque, 'Generating hard satisfiability problems', *Artif. Intell.*, **81**(1-2), 17–29, (1996).

[14] Olaf Steinmann, Antje Strohmaier, and Thomas Stutzle, 'Tabu search vs. random walk', in *KI - Kunstliche Intelligenz*, pp. 337–348, (1997).

[15] John Thornton, Duc Nghia Pham, Stuart Bain, and Valnir Ferreira Jr., 'Additive versus multiplicative clause weighting for sat', in *AAAI*, pp. 191–196, (2004).

[16] Dave A. D. Tompkins and Holger H. Hoos. Novelty$^+$ and adaptive novelty$^+$. SAT 2004 Competition Booklet, 2004. (solver description).

[17] Dave A. D. Tompkins and Holger H. Hoos, 'Warped landscapes and random acts of sat solving', in *Proc. of the Eighth International Symposium on Artificial Intelligence and Mathematics (AIMA-04)*, (2004).

[18] F. Hutter D.A.D. Tompkins and H.H. Hoos, 'Scaling and probabilistic smoothing: Efficient dynamic local search for sat', in *Proc. Principles and Practice of Constraint Programming - CP 2002 : 8th International Conference, CP 2002, Ithaca, NY, USA, September 9-13*, volume 2470 of *LNCS*, pp. 233–248. Springer Verlag, (2002).