LinSBPS

Emir Demirović and Peter J. Stuckey
Department of Computing and Information Systems
University of Melbourne
Australia
(emir.demirovic ∨ pstuckey) @ unimelb.edu.au

I. Introduction

The solver was created with the intention to study the effectiveness of local search inspired techniques for maxSAT. This is a long-term goal where the aim is to develop algorithms that use techniques similar to those of metaheuristics, in particular large neighbourhood search, but within a complete algorithmic setting. Thus, the overall objective would be to improve the *anytime* performance of solvers, which is especially important for large-scale problems where optimality guarantees seem impractical.

II. The Algorithm

We start with the linear MaxSAT algorithm [1]. It computes the optimal solution to a maxSAT problem by repeatedly solving a series of SAT problems, each time adding constraints that force the new solution to be better than than the previously computed one. This algorithm, implemented in Open-WBO [2] with Glucose [3] as the backend solver, was the best solver for the *60 seconds unweighted incomplete track* in the last maxSAT evaluation 2017. However, it was outperformed in the same category with 300 seconds and did not provide competitive solutions for many benchmarks in the weighted incomplete track. The internal SAT solveris a complete backtracking algorithm: it selects a variable, assigns it a truth value, and then either backtracks if a conflict is found or recursively repeats the procedure.

Our approach uses the linear MaxSAT algorithm augmented with two important components: solution-based phase saving and varying resolution technique, where we start considering the problem in *low resolution* and with time increase the resolution.

A. Solution-Based Phase Saving

The variable selection process partially mimics strategies used in local search algorithms: it selects a variable that was frequently involved in recent conflicts (*high activity*, the VSIDS scheme [4]). However, the truth value assignment procedure does not: it is based on *phase saving*, meaning it assigns the value used most recently for the variable. While phase saving is effective for pure SAT problems, *solution-based phase saving* has proven to be more efficient for optimisation [5], where the assignment is based on the best solution found so far. If the previous search was in a space where no better solution exists, time is effectively wasted with standard phase saving. Solution phase saving avoids this by searching around the best solution found. This is reminiscent of local search, as the algorithm is directed *near* the best solution. It can also be seen as a kind of Large Neighbourhood Search [6]. Indeed, assigning values to a set of variables based on the current best soluton and optimising for the remaining variables is a common strategy in metaheuristic algorithms and has been used for decades. Such a technique is particularly relevant for the incomplete track in the maxSAT competition, where solvers are expected to deliver high quality solution within tight time budgets.

To boost its performance, we incorporated solution-based phase saving in the linear algorithm. Solution-based phase saving is not widely used in MaxSAT solving. It is used by WPM3 [7] in a core-guided approach. However, we argue that the technique is more natural for a linear algorithm. As noted, the basic idea has been used in metaheuristic algorithms and even in MaxSAT solving [5] [7], but the position of the linear algorithm with solution-based phase saving among modern MaxSAT solvers is not clear. Thus, we implemented solution-based phase saving in Open-WBO [2] and evaluated its performance using benchmarks from the recent maxSAT evaluation 2017 and the international timetabling competition 2011. We do not present the results of our study in this short paper, but we do note that it provided an improvement over the baseline linear algorithm. In our recent CP paper [8], we studied solution-based phase saving for constraint programming solvers and its relation to *automated large neighbourhood search*. For CP, it provides substantial improvements. We note that we have investigated other phase saving variants, but as of now, the results remain inconclusive.

B. Varying Resolution Approach

While solution-based phase saving does provide improvements, especially for certain classes of problems, it cannot be used effectively for a large set of the MaxSAT competition benchmarks. The reason is that the linear algorithm relies on encoding a single large cardinality constraint, which is directly dependant of the magnitude of the sum of the weights of soft clauses. As the sum grows, in the general case, so does the number of clauses and auxiliary variables that are needed to encode the cardinality constraint. Thus, the memory requirements can be significant. This has a direct impact on the performance of the linear algorithm and it some cases it completely dominates the solver. We note that this is not necessarily the case for core-guided approaches, which for a large part are uneffected by the magnitude of the weights. Hence, we developed a simplification strategy where we initially consider the problem in *low resolution* where the weights of the MaxSAT problem are divided by a large value. After the simplified problem is solved optimally, the resolution of the problem is increased i.e. the division value is lowered. This continues iteratively until the full original problem is solved. Therefore, the technique is theoretically complete, but in practice for the benchmarks from the last MaxSAT evaluation and the short time limits, only one or two resolutions are typically considered. We note that solution-based phase saving is used during the algorithm, as well as in between resolutions. With this technique, intuitively, the most important constraints are dealt with in the beginning and with executation time other increasing important constraints are added the clause database, resembling local search style methods. It is related to the lexicographical optimisation approach for MaxSAT [9].

The main advantage is that the cardinality constraint that needs to be encoded is orders of magnitude smaller than from the original problem, offering substantial speed-ups. However, the varying resolution approach comes at the price of precision, as an optimum solution for the low-resolution problem does not necessarily correspond to the optimum for the higher resolutions and vice versa. Moreover, given two models for the low-resolution problem and their cost, it is not possible to determine which one of them is better based on their cost without consider the complete original problem. The tendency, heuristically speaking, is that better solutions to the low-resolution problem correlate with better solutions to the original problem.

III. Conclusion

We presented LinSBPS, the algorithm we submitted for the MaxSAT Evaluation 2018. It uses a linear MaxSAT algorithm coupled with solution-based phase saving and a varying resolution approach. Our experimental results have shown that significant improvements could be achieved when compared with *maxroster*, one of the top performing solvers from the incomplete track last year. However, more detailed experimental results, such as those provided by the MaxSAT competition, are required to draw stronger conclusions.

References

[1] D. Le Berre and A. Parrain, "The Sat4j library, release 2.2 system description," *Journal on Satisfiability, Boolean Modeling and Computation*, vol. 7, pp. 59–64, 2010.

[2] R. Martins, V. Manquinho, and I. Lynce, "Open-WBO: a modular maxSAT solver," in *Proceedings of SAT-14*, pp. 438–445.

[3] G. Audemard and L. Simon, "Predicting learnt clauses quality in modern sat solvers," in *Proceedings of IJCAI'09*.

[4] M. W. Moskewicz, C. F. Madigan, Y. Zhao, L. Zhang, and S. Malik, "Chaff: Engineering an efficient SAT solver," in *Proceedings of DAC'01*, pp. 530–535.

[5] I. Abío Roig, "Solving hard industrial combinatorial problems with SAT," *Ph.D. thesis, Technical University of Catalonia (UPC)*, 2013.

[6] P. Shaw, "Using constraint programming and local search methods to solve vehicle routing problems," in *Proceedings of CP'98* (M. Maher and J.-F. Puget, eds.), pp. 417–431, Springer.

[7] C. Ansótegui, F. Didier, and J. Gabàs, "Exploiting the structure of unsatisfiable cores in maxSAT," in *Proceedings of IJCAI-15*, pp. 283–289.

[8] E. Demirović and P. J. Stuckey, "Solution-based phase saving and large neighbourhood search," in *to appear in the proceedings of CP'18'*.

[9] J. Marques-Silva, J. Argelich, A. Graça, and I. Lynce, "Boolean lexicographic optimization: algorithms & applications," *Annals of Mathematics and Artificial Intelligence*, vol. 62, no. 3-4, pp. 317–343, 2011.