

ParLS-PBO: A Parallel Local Search Solver for Pseudo Boolean Optimization

Zhihan Chen ✉ 

Key Laboratory of System Software (Chinese Academy of Sciences) and State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, Beijing, China
School of Computer Science and Technology, University of Chinese Academy of Sciences, Beijing, China

Peng Lin ✉ 

Key Laboratory of System Software (Chinese Academy of Sciences) and State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, Beijing, China
School of Computer Science and Technology, University of Chinese Academy of Sciences, Beijing, China

Hao Hu ✉ 

Key Laboratory of System Software (Chinese Academy of Sciences) and State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, Beijing, China

Shaowei Cai¹ ✉ 

Key Laboratory of System Software (Chinese Academy of Sciences) and State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, Beijing, China
School of Computer Science and Technology, University of Chinese Academy of Sciences, Beijing, China

Abstract

As a broadly applied technique in numerous optimization problems, recently, local search has been employed to solve Pseudo-Boolean Optimization (PBO) problem. A representative local search solver for PBO is *LS-PBO*. In this paper, firstly, we improve *LS-PBO* by a dynamic scoring mechanism, which dynamically strikes a balance between score on hard constraints and score on the objective function.

Moreover, on top of this improved *LS-PBO*, we develop the first parallel local search PBO solver. The main idea is to share good solutions among different threads to guide the search, by maintaining a pool of feasible solutions. For evaluating solutions when updating the pool, we propose a function that considers both the solution quality and the diversity of the pool. Furthermore, we calculate the polarity density in the pool to enhance the scoring function of local search. Our empirical experiments show clear benefits of the proposed parallel approach, making it competitive with the parallel version of the famous commercial solver *Gurobi*.

2012 ACM Subject Classification Computing methodologies → Parallel algorithms; Theory of computation → Randomized local search

Keywords and phrases Pseudo-Boolean Optimization, Parallel Solving, Local Search, Scoring Function, Solution Pool

Digital Object Identifier 10.4230/LIPIcs.CP.2024.5

Supplementary Material *Software*: <https://github.com/shaowei-cai-group/ParLS-PBO.git>
archived at `swh:1:dir:7aa353d93bbe537f550a8c0e4bea165b7cc2f136`

Funding This work is supported by National Key R&D Program of China (2023YFA1009500).

¹ Corresponding author



1 Introduction

With the recent impressive progress in high-performance Boolean Satisfiability (SAT) and Maximum Boolean Satisfiability (MaxSAT) solvers, increasing real-world problems are solved with the Conjunctive Normal Form (CNF) encoding. However, in practice, CNF is ineffective in dealing with cardinality constraints, resulting in its size growing dramatically [3]. As a rich subject in various fields, Pseudo-Boolean Optimization (PBO) provides a better formalization than CNF in expressive power, with the use of Linear Pseudo-Boolean (LPB) constraints. Meanwhile, LPB constraints stay close to CNF and can benefit from advances in SAT solving [33]. The PBO problem is to find an assignment satisfying all LPB constraints that maximizes the objective function given.

Briefly, there are three categories of complete algorithms to solve the PBO problem. The first one is the linear search, which extends the PB solver by adding a constraint enforcing to find a better solution (in terms of the objective function value) when finding a solution satisfying all constraints [2]. Several well-known PBO solvers are based on this idea, including *Sat4j* [22], *RoundingSAT* [14], and *HYBRID* [12]. The second one is Branch-and-Bound, which focuses on the techniques to estimate the lower bounds of the objective value, as the search can be pruned whenever the lower bound is greater than or equal to the upper bound. The symbolic techniques to determine the lower bounds include Maximum Hitting Set [11] and Linear Programming Relaxation [26]. The third one is to call the SAT solvers after encoding PB constraints into the CNF formula, such as *MINISAT+* [13] and *OpenWBO* [31]. Moreover, mixed-integer programming (MIP) solvers can be directly applied to solve the PBO problem, as PB constraints can be treated as 0-1 linear constraints, representative solvers include *SCIP* [5] and *Gurobi* [17].

Complete algorithms often suffer from the scalability issue, which motivates the development of incomplete algorithms. A typical incomplete approach is local search, which has been successfully used in many problems, including SAT [25, 1, 6], MaxSAT [23, 9], etc. Nevertheless, the literature on local search algorithms to solve PBO problem is quite limited. The first local search-based PBO solver was proposed in [24], called *LS-PBO*. This local search solver introduced a constraint weighting scheme and a scoring function considering both hard and soft constraints to select Boolean variables to flip. Later, *LS-PBO* was improved by using a unit propagation-based method to produce better initial assignments [19], resulting in the *DeciLS-PBO* solver. Very recently, on top of *LS-PBO*, Chu et. al. developed *NuPBO* [10], which established the latest state-of-the-art local search based PBO solving. Additionally, Iser et. al. proposed an oracle-based local search approach in the context of PBO [18], which outperforms on various benchmark domains clearly the recent pure stochastic local search approach.

Recently, with the evolution of multi-core processors, parallel solving received growing interest. The SAT competition² set up a parallel track from 2009, while Satisfiability Modulo Theories (SMT) competition³ introduces parallel tracks in 2021. In short, parallel algorithms contain two major directions. The first one is based on the concept of divide-and-conquer, which divides the problem into several sub-problems, and each thread solves sub-problems. For example, Treengeling [15] is a representative SAT solver of this kind. Meanwhile, commercial solvers, such as *CPLEX*⁴ and *Gurobi* [17] also implement their parallel versions via this approach. The other parallel approach is to integrate different solvers, including a solver with

² <http://www.satcompetition.org/>

³ <https://smt-comp.github.io/>

⁴ <http://www.cplex.com/>

different configurations, and each thread runs a solver. This approach is commonly known as portfolio, which is simple but effective. The portfolio-based parallel SAT solvers, such as *PRS* [8], *P-mcomsps* [16], and *Pakis* [36], dominate the parallel track of SAT competitions in recent years. The parallel MaxSAT solvers [29, 30] based on the portfolio method also demonstrate a strong ability to efficiently solve a large number of problem instances due to the use of complementary search strategies and sharing learned clauses between threads.

In this paper, at first, we improve the typical *LS-PBO* solver by introducing a new dynamic scoring mechanism, which can somehow avoid the local optimum situation even after flipping thousands of variables. This leads to an improved algorithm called *DLS-PBO*. Then, based on the *DLS-PBO*, we develop *ParLS-PBO*, to the best of our knowledge, the first parallel local search-based PBO solver. Our parallel solver runs different local search procedures in the worker threads and maintains a solution pool, which collects good feasible solutions from the working threads and in turn can be used to guide the local search procedures. During the search process, the solution pool is updated by adding new solutions and removing solutions from it. To update the solution pool, we propose a function to measure the feasible solutions found by local search, which considers both the objective value of the solution and the diversity of the solutions in the pool.

The most important part of our parallel solver is how to use the solution pool to help the local search. In this work, this is done in two ways. Firstly, the solutions in the pool can be directly used to help local search when it stagnates for a long time. Specifically, in such a situation, a local search process of a thread restarts from a good feasible solution picked from the pool. Secondly, we calculate the polarity density (the proportions of being 1 and 0) for each variable once a solution is added into the solution pool. This polarity density information is used to enhance the scoring function of local search when picking the variable to flip in each step. The intuition is when a certain polarity (either 1 or 0) of a variable occurs in most high-quality solutions, it brings preference to assign the variable to that polarity.

We carry out experiments to evaluate our algorithms *DLS-PBO* and *ParLS-PBO* on both real-world applications encoded benchmark and standard benchmarks, compared with state-of-the-art solvers including *LS-PBO* [24], *DeciLS-PBO* [19], *NuPBO* [10], *SCIP* [5], *HYBRID* [12], *PBO-IHS* [35], *Gurobi* [17]), and *FiberSCIP* [34]. Our results show that our parallel solver has significantly better performance than all sequential solvers, and competes well with the parallel versions of *Gurobi*. Furthermore, *ParLS-PBO* shows good scalability up to 32 threads as its performance improves with the number of threads.

The remainder of this paper is structured as follows. Section 2 introduces preliminary knowledge. Section 3 analyzes the weakness of *LS-PBO* and introduces an improved solver *DLS-PBO*. Section 4 presents the proposed parallel solver *ParLS-PBO*. Experimental studies are presented in Section 5. Finally, we give some concluding remarks in Section 6.

2 Preliminaries

2.1 Pseudo-Boolean Optimization

A Boolean variable x_i can take only two values *false* and *true*, or equivalently $\{0, 1\}$. A literal l_i is either a variable x_i or its negation $\neg x_i$. Given a set of n Boolean variables $\{x_1, \dots, x_n\}$, a *linear pseudo-Boolean constraint* (LPB constraint) is formed as follows:

$$\sum_{i=1}^n a_i \cdot l_i \triangleright b, \quad a_i, b \in \mathbb{Z}, \quad \triangleright \in \{=, \leq, <, \geq, >\}$$

where a_i is the coefficient for literal l_i , b is called as the degree of the constraint, and \triangleright is one of the classical relational operators. With a given assignment or partial assignment, the constraint is satisfied if its left and right terms satisfy the relational operator. Otherwise, it is unsatisfied.

Moreover, replacing all literals x_i (respectively $\neg x_i$) with negative coefficients with $1 - \neg x_i$ (respectively $1 - x_i$), a LPB constraint can be normalized into the following form [33]:

$$\sum_{i=1}^n a_i \cdot l_i \geq b, \quad a_i, b \in \mathbb{N}_0^+$$

Given a conjunction of LPB constraints, *Pseudo-Boolean Solving* (PBS) problem is a decision problem to find an assignment such that all constraints are satisfied. *Pseudo-Boolean Optimization* (PBO) problem is an optimized version of the PBS problem, aiming to find an assignment satisfying all constraints with the minimal value of a given objective function. In this paper, we focus on the PBO problem consisting of a conjunction of LPB constraints and a linear objective function. Therefore, a PBO instance subjecting to m LPB constraints has the following form:

$$\begin{aligned} & \min_{\{x_1, \dots, x_n\}} \sum_{i=1}^n c_i \cdot l_i, \quad c_i \in \mathbb{Z} \\ & \text{subject to: } \bigwedge_{j=1}^m \sum_{i=1}^n a_{ji} \cdot l_i \geq b_j, \quad a_{ji}, b_j \in \mathbb{N}_0^+ \end{aligned}$$

where c_i is the objective coefficient for literal l_i .

2.2 A Review of *LS-PBO* Solver

LS-PBO is a representative local search solver for PBO, and serves as the basis of other local search PBO solvers. Briefly, it contains two main ideas: a *Constraint Weighting Scheme* and *Scoring Functions* for guiding the search process.

To solve a standard PBO instance, *LS-PBO* proposed a soft *objective constraint*: $\sum_{i=1}^n c_i \cdot l_i < obj^*$, where obj^* indicates the objective value of the best solution in the current run, and other constraints are set as hard. *LS-PBO* uses a weighting technique to increase the weights of falsified constraints, so that the search process is biased toward satisfying them. Specifically, it used dynamic weights (denoted as $w(\cdot)$) to help the search avoid stuck in the local optimum, while increasing the weights of hard constraints to find feasible solutions, and the weight of objective constraint to find better solutions.

Besides, scoring functions are essential in local search algorithms to guide the search process, which typically measures the benefits of flipping a Boolean variable. In *LS-PBO*, the score of flipping a variable x (denoted as $score(x)$) was defined as follow:

$$score(x) = hscore(x) + oscore(x) \tag{1}$$

where $hscore(x)$ indicates the decrease of the total penalty of falsified hard constraints caused by flipping x , and $oscore(x)$ indicates the decrease of the penalty of the objective constraint caused by flipping x . In detail, the penalty of falsifying a hard constraint hc was defined as $w(hc) \cdot \max(0, b - \sum_{i=1}^n a_i \cdot l_i)$, and the penalty for the objective constraint oc was defined as $w(oc) \cdot \sum_{i=1}^n c_i \cdot l_i$.

3 Improving *LS-PBO* Solver with Dynamic Scoring Mechanism

As introduced in the preliminary, the score of a flipping variable x ($score(x)$) in *LS-PBO* is presented as Equation 1. The algorithm selects the variable with the highest positive score, indicating the biggest decrease in the penalty of hard constraints and objective constraint. A drawback of *LS-PBO* is the lack of dynamic adjustments to the ratio of the soft and hard constraints. If a feasible solution cannot be found within a certain period of time, the search mechanism should adaptively prioritize finding feasible solutions, thereby increasing the ratio attributed to the hard constraints. Conversely, if feasible solutions have been frequently found recently, then it would be beneficial to increase the ratio of the soft constraints to guide the search towards discovering better solutions.

To resolve this drawback, we introduce a new *dynamic scoring function*, denoted as $score^*(x)$, to adjust the significance of $oscore(x)$ for every given K steps (K is a parameter), which is defined as follows:

$$score^*(x) = hscore(x) + p \cdot oscore(x) \quad (2)$$

where p is a dynamic ratio initially set as 1. It would be decreased as p/inc (where $inc > 1$) if no feasible solution is found during the recent K steps, to guide the search towards a feasible solution. Otherwise, respectively, it would be increased as $p \cdot inc$ when a feasible solution is found within the recent K steps, to guide the search process for a better solution.

► **Example 1.** Considering a PBO instance:

$$\begin{aligned} \min_{\{x_1, x_2, x_3\}} \quad & 10 \cdot x_1 + 20 \cdot x_2 + 30 \cdot x_3 \\ \text{subject to:} \quad & 2 \cdot x_1 + 3 \cdot x_2 + 4 \cdot x_3 \geq 5 \end{aligned}$$

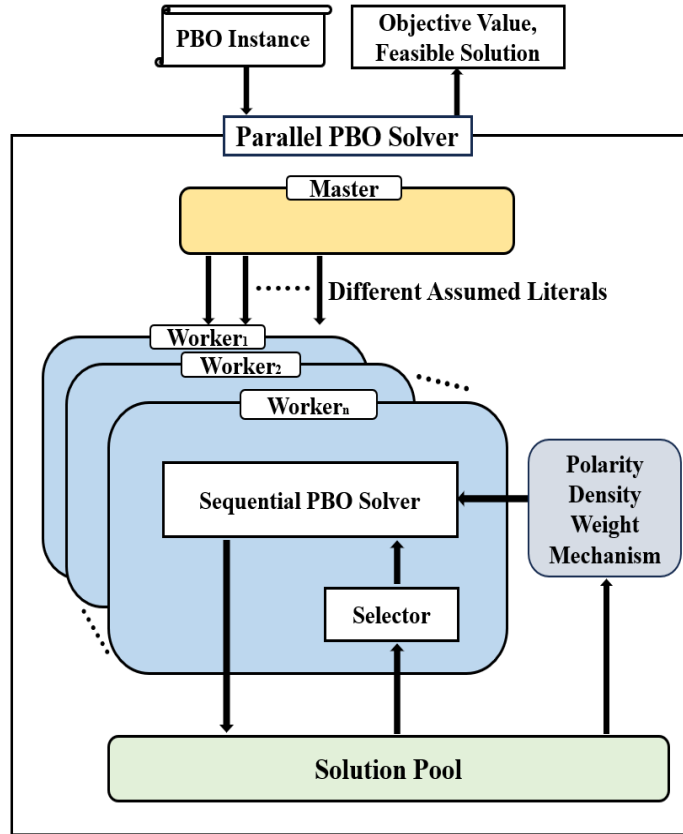
and suppose current weights $w(hc)$ and $w(oc)$ are 2 and 1. For the given assignment ($x_1 = 1, x_2 = 0, x_3 = 0$), the corresponding $hscore(\cdot)$ and $oscore(\cdot)$ are as follows:

\cdot	x_1	x_2	x_3
$hscore(\cdot)$	-4	6	6
$oscore(\cdot)$	10	-20	-30

Consider the following two situations:

- If feasible solutions are found frequently in recent period, the value of p will gradually increase, guiding the search to lower the cost of the objective constraint. Suppose the current value of p is 2, then $score^*(x_1) = 16$, $score^*(x_2) = -34$, $score^*(x_3) = -54$. In this case, x_1 will be picked and flipped, resulting in a decrease of 10 in the cost of the objective constraint. (even if it is not a feasible solution.)
- If the algorithm has not visited feasible solutions for a period (K steps), the value of p will gradually decrease, guiding the search to find feasible solutions. Suppose the current value of p is 0.1. then $score^*(x_1) = -3$, $score^*(x_2) = 4$, $score^*(x_3) = 3$. In this case, x_2 will be picked and flipped, resulting in a feasible solution.

We denote the improved version of *LS-PBO* solver with dynamic scoring mechanism as *DLS-PBO*.



■ **Figure 1** Architecture of the *ParLS-PBO*.

4 Parallel Local Search Solver for Pseudo-Boolean Optimization

In this section, we propose a parallel local search solver *ParLS-PBO*. The architecture of *ParLS-PBO* is shown in Figure 1, which consists of two major contributions: *Solution Pool* and *Polarity Density Weight*. We first describe the global framework of *ParLS-PBO*, then we present the contributions in detail separately.

4.1 Framework of *ParLS-PBO* Solver

As a portfolio-based local search PBO solver, *ParLS-PBO* contains a *master* thread and multiple *worker* threads. The master thread reads the input PBO instance, and then produces different initial partial assignments via *literal assume* technique for worker threads; finally, when the time limit is reached, it outputs the best solution returned from all worker threads. In detail, supposing there are T worker threads, the master thread selects $\lceil \frac{T}{2} \rceil$ random variables. Then, for each variable selected x_i , it generates a positive literal x_i and a negative literal $\neg x_i$. Therefore, it generates T (or $T + 1$ if T is odd) different assumed literals in total for worker threads.

Each worker thread receives an assumed literal ℓ (either x_i or $\neg x_i$) and applies the *unit propagation* [28] technique to simplify the formula. Note that a solution found by local search for such a formula can be directly transformed to a solution for the original PBO instance, by adding the assumed literal as the value for the corresponding variable, and adding the value for reduced variables via unit propagation if any. Then the worker thread launches a local

search solver to solve the PBO instance. In default, the worker thread launches *DLS-PBO*. To bridge different worker threads, we propose a *Solution Pool* to share high-quality feasible solutions found from different worker threads. When the search process of a worker thread is blocked in the local optimum after flipping a certain number of variables, it attempts to restart with a high-quality feasible solution in the solution pool. Furthermore, we introduce the concept of *Polarity Density Weight* with the intuition of preference of certain polarity of a variable if it occurs in most high-quality solutions.

4.2 Maintaining the Solution Pool

The solution pool aims to collect good feasible solutions, preferring those with more differences. To this end, we consider a mixed quality rating function $r_{mix}(\cdot)$ by measuring two terms together: the quality in objective value, and the diversity. For a feasible solution S , the rating function $r_{mix}(S)$ w.r.t. a solution pool is defined as follows:

$$r_{mix}(S) = rank_{obj}(S) \cdot p^* + rank_{div}(S) \cdot (1 - p^*) \quad (3)$$

where $rank_{obj}(\cdot)$ and $rank_{div}(\cdot)$ represent the ranking of S in the solution pool in objective value and diversity value. Specifically, for the objective value, the solution with the minimal objective value is considered as the best solution, hence its $rank_{obj}(\cdot)$ value is 1. While for the diversity value, the solution with the maximum diversity value is considered the best, thus its $rank_{div}(\cdot)$ value is 1. p^* is a penalty parameter within $[0, 1]$ to adjust the significance of the objective value term and diversity term.

The difference between two solutions is measured as the sum of the number of different polarities, and the diversity value of a solution S w.r.t. The solution pool is measured as the sum of differences between S and all other solutions in the solution pool. Formally,

$$div(S) = \sum_{S' \in \mathcal{P}} Hamming(S, S')$$

When a worker thread finds a new feasible solution S , if the solution pool is not full, then just add it. Otherwise, S replaces the worst one (the solution with the biggest $r_{mix}(\cdot)$ value).

We note that the $r_{mix}(S)$ function in this work resembles a previous population management strategy [7]. We focus on the ranking rather than the value, which can be seen as a normalization.

4.3 Using the Solution Pool to Guide the Search

In this subsection, we discussed how the solution pool guides each worker, including replacing solutions with better solutions from the solution pool when a worker being trapped, and utilizing the variable polarity preference in the solution pool to influence the selection of variables to flip during the search process.

4.3.1 Solution Sharing Strategy

When a worker thread fails to find a better feasible solution for a while, which means that it may be trapped in a local optimum, it selects a feasible solution with a smaller objective value from the solution pool and replaces the current one.

In practice, each worker thread preserves the current best feasible solution (denoted as S^*) as well as the corresponding objective value obj^* . When the search process fails to find a better solution after R steps, it picks a solution from the solution pool as a new starting point.

To prevent excessive overlap of search spaces among various threads, we employ a probability-based method to select solutions in the pool, rather than directly choosing the best solution in the pool. Specifically, let $\{S_1, \dots, S_k\}$ denotes the set of feasible solutions in the solution pool with objective values not bigger than obj^* (the set will not be empty, as it at least contains S^*), and Δ_i denotes the difference between the objective value of S_i and obj^* . Then the probability of selecting S_i is $\Delta_i / \sum_{j=1}^k \Delta_j$.

4.3.2 Polarity Density Weight

Besides using the solutions in the pool to guide the search process directly when it gets stuck, we propose a deeper guiding method, which utilizes a piece of valuable hidden information in the solution pool – the occurrence of polarities (0 or 1) of variables. To measure the effect of this kind of information, we propose the concept of *polarity density weight* for a variable x , denoted as $w_{pd}(x)$, which reflects the preference of certain polarity of x appearing in high-quality solutions.

In detail, for a variable x , $w_{pd}(x)$ is initialized as 1. Once a high-quality feasible solution S is added into the solution pool, $w_{pd}(x)$ will add (respectively, minus) a step value β when positive (respectively, negative) polarity of x appears in S . In fact, via this updating mechanism, the higher (respectively, lower) value in $w_{pd}(x)$ indicates the higher preference of positive (respectively, negative) polarity for x in high-quality solutions. To limit the influence of $w_{pd}(x)$ and avoid possible calculation problems in negative values, we restrict $w_{pd}(x)$ into an interval of $[1 - \epsilon, 1 + \epsilon]$, where ϵ scales the bound. Therefore, the $w_{pd}(x)$ is updated as follows:

$$w_{pd}(x) = \begin{cases} \max(w_{pd}(x) - \beta, 1 - \epsilon), & \text{if } x = 0 \text{ in } S \\ \min(w_{pd}(x) + \beta, 1 + \epsilon), & \text{if } x = 1 \text{ in } S \end{cases} \quad (4)$$

The polarity density weight is used to enhance the scoring function of picking a variable to flip during the search process. The resulting enhanced scoring function, denoted as $score^{**}(x)$, is defined as follows:

$$score^{**}(x) = \begin{cases} score^*(x) \cdot w_{pd}(x), & \text{if } x = 0 \text{ in } S_{cur} \\ score^*(x) / w_{pd}(x), & \text{if } x = 1 \text{ in } S_{cur} \end{cases} \quad (5)$$

where S_{cur} is the current assignment maintained by the local search process.

The multiplication of polarity density weight influences the flip of a variable x from 0 to 1, as it increases the combined score if the preference of positive polarity exists ($w_{pd}(x) > 1$) to guide the search process to realise the flip. Respectively, in reverse, the division of polarity density weight influences the flip from 1 to 0.

► **Example 2.** Continuing with Example 1, Suppose that most of the solutions that entered the solution pool have the assignment $(x_1 = 1, x_2 = 1, x_3 = 0)$, resulting in $w_{pd}(x_1) = 1.1, w_{pd}(x_2) = 1.1, w_{pd}(x_3) = 0.9$.

For the given assignment $(x_1 = 1, x_2 = 0, x_3 = 0)$ and $p = 1$, the corresponding $score^*(\cdot)$ can be calculated as: $score^*(x_1) = 6, score^*(x_2) = -14, score^*(x_3) = -24$.

Then $score^{**}(x_1) = 6 \div 1.1, score^{**}(x_2) = (-14) \times 1.1, score^{**}(x_3) = (-24) \times 0.9$.

5 Experiments

The experiments are organized as three parts. At first, we focus on comparing *DLS-PBO*, *ParLS-PBO* with state-of-the-art solvers including commercial solvers. Secondly, we analyze the effectiveness of the strategies to guide the search via the solution pool in *ParLS-PBO*. Finally, we present the tendency in performance of *ParLS-PBO* with the increase of the number of threads. Source code and detailed results are made publicly available on GitHub⁵.

5.1 Benchmark

- **Real-World:** Three real-world application problems, which are presented in the literature [24], including the Minimum-Width Confidence Band Problem [4]⁶ (24 instances), the Seating Arrangements Problem [32] (21 instances), the Wireless Sensor Network Optimization Problem [20, 21] (18 instances).
- **MIPLIB:** All satisfiable 0-1 integer programs from the MIPLIB 2017 library and earlier MIPLIB releases⁷, which contains 252 instances, provided in the literature [35].
- **PB16:** The OPT-SMALL-INT benchmark from the most recent Pseudo-Boolean Competition 2016⁸. We filter out the duplicated instances that appear in both MIPLIB and PB16, resulting in 1524 instances in the final. PB16 contains different problem categories. We select those representatives (containing more than 30 instances) categories for finer-grained experimental analysis.

5.2 Candidate Methods to Compare

In the sequential track, we compare *DLS-PBO* with 7 state-of-the-art sequential PBO solvers, including 3 local search-based solvers: *LS-PBO*, *DeciLS-PBO* and *NuPBO*, 3 complete non-commercial solvers: *HYBRID*, *PBO-IHS*, and *SCIP* and the commercial solver *Gurobi* (both complete and heuristic versions).

In the parallel track, we compare *ParLS-PBO* with the academic solver *FiberSCIP*, and the parallel version of the commercial solver *Gurobi*.

- *LS-PBO* [24]: the state-of-the-art SLS algorithm for solving PBO⁹.
- *DeciLS-PBO* [19]: a recent SLS algorithm based on LS-PBO¹⁰.
- *NuPBO* [10]: a recent SLS algorithm based on LS-PBO, which established the latest state-of-the-art local search based PBO solving¹¹.
- *HYBRID* [12]: a recent core-guided PBO solver building upon RoundingSAT [14]¹².
- *PBO-IHS* [35]: a recent IHS PBO solver building upon RoundingSAT¹³.
- *Gurobi* [17]: one of the most powerful commercial MIP solvers. We use both its complete and heuristic versions¹⁴.

⁵ <https://github.com/shaowei-cai-group/ParLS-PBO.git>

⁶ <http://physionet.org/physiobank/database/mitdb/>

⁷ <https://zenodo.org/record/3870965>

⁸ <http://www.cril.univ-artois.fr/PB16/bench/PB16-used.tar>

⁹ <https://lcs.ios.ac.cn/~caisw/Resource/LS-PBO/>

¹⁰ <https://github.com/jiangluyu1998/DeciLS-PBO> (commit number: 3dce881)

¹¹ <https://github.com/filyouzicha/NuPBO> (commit number: 821d901)

¹² <https://zenodo.org/record/4043124> (version 2)

¹³ <https://bitbucket.org/coreo-group/pbo-ihs-solver> (version 1.1)

¹⁴ <https://www.gurobi.com/solutions/gurobi-optimizer> (version 10.0.0)

- *SCIP* [5]: one of the fastest non-commercial solvers for MIP (the latest version 8.0.1)¹⁵,
- *FiberSCIP* [34]: a parallel non-commercial MIP solvers based on SCIP (the latest version 1.0.0)¹⁶.

We download the latest version of all candidate methods to compare from their published links. In all experiments, we always use their default parameter settings.

5.3 Experimental Settings

DLS-PBO and *ParLS-PBO* are implemented in C++, and compiled with g++ (version 9.2.0) using the option '-O2'. All experiments are carried out on a cluster with two AMD EPYC 7763 CPUs @ 2.45Ghz of 128 physical cores and 1TB memory running the operating system Ubuntu 20.04 LTS (64bit).

As with the previous research on PBO solvers [24, 19], we set the time limit for each run as 300 and 3600 seconds. For each sequential randomized solver, we run 10 times for each instance with different seeds from $\{0, 1, \dots, 9\}$, and select the median of the 10 runs as the final result. Without making any additional claims, the number of CPU cores that can be used for parallel solvers is set as 32.

For parameter tuning, we employed Sequential Model-based Algorithm Configuration (SMAC) [27], conducting the tuning on 300 instances randomly selected from all the benchmarks, with a time limit set to 300 seconds. The parameter values obtained after tuning are listed in Table 1¹⁷.

■ **Table 1** The parameter settings of our solvers.

Parameter	K	R	inc	$poolsize$	p^*	β	ϵ
Value	566024	86295	1.15	18	0.58	0.03	0.144

Referring to the MaxSAT competition and previous research on PBO, we use 2 metrics to evaluate the performance of each solver:

- $\#win$: the number of instances that a solver finds the best solution among all solutions output by tested solvers (i.e., the number of winning instances).
- avg_{sc^*} : Since 2017, in the incomplete track of recent MaxSAT Evaluations, the performance of various solvers is measured by competition scores. For an instance and a solver given, the competition score sc is defined as $(1 + cost_{best}) / (1 + cost_s)$, where $cost_{best}$ represents the objective value of the best solution found among all solvers, $cost_s$ represents the objective value of the solution found by the given solver. However, in PBO problem, the objective value of a solution may be negative, leading to an incorrect calculation of sc . To address this issue, we modify slightly the definition of competition score:

$$sc^* = \frac{1 + cost_{best} + \sum_{c_i < 0} |c_i|}{1 + cost_s + \sum_{c_i < 0} |c_i|}$$

Adding all negative objective coefficients ensures the competition score of each instance is normalized in $[0, 1]$. We use avg_{sc^*} to denote the average competition score of a solver.

¹⁵ <https://www.scipopt.org/index.php#download> (version 8.0.1)

¹⁶ <https://ug.zib.de/index.php#download> (version 1.0.0)

¹⁷ In fact, our solver is not sensitive to the parameter configurations. For example, a simple configuration ($K=100000$, $R=100000$, $inc=1.1$, $poolsize=10$, $\beta=0.1$, $\epsilon=0.15$) leads to a performance close to the one in Table 1, with a gap of avg_{sc^*} less than 1%.

We do not use average time as a metric because our primary focus is on the quality of the solution. If the quality of the solutions found is different, then the comparison based on run time would be misleading.

5.4 Performance Evaluations

5.4.1 The Sequential Track

We first compare *DLS-PBO* with *LS-PBO*, and the results are shown in Table 2. *DLS-PBO* significantly improves *LS-PBO* in terms of both $\#win$ and avg_{sc^*} on all the benchmarks.

Further, we evaluate *DLS-PBO* with other PBO solvers, as well as integer programming solvers. The results (Table 3) indicate that *NuPBO* performs best for the Real-world benchmark, while *Gurobi* is the best on MIPLIB and PB16 benchmarks. *DLS-PBO* cannot rival these two solvers, yet it is better than other PBO solvers. We note that the emphasis of this work is to develop an effective parallel method for PBO solvers. We choose *LS-PBO* as the baseline as it is the typical local search PBO solver (*NuPBO* is also developed on top of it). We simply remedy its drawback to obtain *DLS-PBO*, and do not perform other modifications. *NuPBO* was published very recently, and we believe our parallel method can be applied to *NuPBO* as well.

5.4.2 The Parallel Track

The comparative results of our parallel solver *ParLS-PBO* with other parallel solvers are shown in Table 4 (We only show $\#win$ due to the space limit). *ParLS-PBO* gives the best performance on all categories of the Real-World benchmark, and 3 categories of the PB16 benchmark, including Kexu, Logic Synthesis, and Prime.

In terms of the Total instances, *ParLS-PBO* outperforms the non-commercial solver *FiberSCIP*, and is competitive with the commercial solver *Gurobi*. Comparing Table 3 and Table 4, it can be found that the gap between *ParLS-PBO* and *Gurobi* (32 threads) is decreasing compared with the gap between *DLS-PBO* and *Gurobi* (1 thread), which indicates the effectiveness of our solver in parallel solving.

We also observe that *ParLS-PBO* outperforms the best sequential PBO solver *NuPBO* on all benchmarks (44 vs. 32, 171 vs. 156, and 1238 vs. 1002). Although this comparison is unfair (and thus we do not report it in the table), it indicates that by parallelization, the performance of PBO solvers can be significantly improved.

5.5 Effectiveness Analysis

This subsection evaluates the effectiveness of the key strategies of *ParLS-PBO*. In Table 5, we compare *ParLS-PBO* with its 2 variants:

- V_1 : to analyze the effectiveness of the solution-pool-based sharing, we modify *ParLS-PBO* by disabling the sharing mechanism and making each thread solve separately.
- V_2 : to analyze the effectiveness of the global score mechanism, we modify *ParLS-PBO* by disabling the global score mechanism and using $score^*(x)$ directly in the local search.

As shown in Table 5, *ParLS-PBO* outperforms other variations, confirming the effectiveness of the strategies.

■ **Table 2** Evaluation between *DLS-PBO* and *LS-PBO*.

Benchmark	#Ins	<i>LS-PBO</i>		<i>DLS-PBO</i>	
		#win	avg _{sc} *	#win	avg _{sc} *
cutoff=300s					
Real-World	63	25	0.976	48	0.996
miplib	252	118	0.777	182	0.836
PB16	1524	711	0.692	1124	0.776
Total	1839	854	0.713	1354	0.792
cutoff=3600s					
Real-World	63	38	0.991	35	0.988
miplib	252	121	0.81	186	0.863
PB16	1524	829	0.753	1189	0.825
Total	1839	988	0.769	1410	0.836

■ **Table 3** Performance evaluation between *DLS-PBO* and sequential SOTA solvers (The results of *DeciLS-PBO* are not presented due to the space limit. In fact, *DeciLS-PBO* is dominated by *NuPBO* and *DLS-PBO*).

Benchmark	#Ins	<i>SCIP</i>	<i>HYBRID</i>	<i>PBO-IHS</i>	<i>NuPBO</i>	<i>Gurobi(Comp.)</i>	<i>Gurobi(Heur.)</i>	<i>DLS-PBO</i>
		#win avg _{sc} *	#win avg _{sc} *	#win avg _{sc} *	#win avg _{sc} *	#win avg _{sc} *	#win avg _{sc} *	#win avg _{sc} *
cutoff=300s								
Real-World	63	0 0.126	3 0.109	2 0.266	46 0.972	4 0.289	4 0.292	29 0.977
MIPLIB	252	88 0.614	53 0.572	81 0.741	116 0.854	152 0.838	165 0.849	101 0.803
PB16	1524	810 0.687	663 0.624	882 0.804	980 0.813	1071 0.84	1072 0.84	842 0.741
cutoff=3600s								
Real-World	63	0 0.171	11 0.494	5 0.401	43 0.997	11 0.34	9 0.38	27 0.974
MIPLIB	252	113 0.675	65 0.697	96 0.789	113 0.859	171 0.893	181 0.895	106 0.831
PB16	1524	906 0.734	729 0.715	939 0.814	1012 0.822	1138 0.86	1144 0.862	940 0.797

5.6 Scalability Analysis

In order to analyze the scalability of *ParLS-PBO*, we choose *Gurobi* (complete version) with 32 threads as the comparison baseline to test the performance gap between different threads of *ParLS-PBO*. We report #win for threads set to {4, 8, 16, 32} compared to baseline. As is shown in Figure 2, in each benchmark, #win is gradually increasing, which verifies the scalability of *ParLS-PBO*.

■ **Table 4** Performance evaluation between *ParLS-PBO* and parallel SOTA solvers.

Benchmark	Category	#ins	<i>FiberSCIP</i>	<i>Gurobi</i>		<i>ParLS-PBO</i>
				Comp.	Heur.	
cutoff=300s						
Real-World	MWCB	24	0	0	0	24
	WSNO	18	0	4	4	18
	SAP	21	0	0	0	21
	Total	63	0	4	4	63
MIPLIB	Total	252	113	190	180	129
PB16	Factor	192	186	186	186	172
	Kexu	40	6	10	7	40
	Logic synthesis	74	71	73	72	73
	Market split	40	12	21	13	5
	Mps	35	30	33	34	23
	Numerical	34	13	18	21	8
	Prime	156	123	128	129	131
	Reduced mps	273	76	145	150	39
	Total	1524	898	1147	1143	1100
cutoff=3600s						
Real-World	MWCB	24	0	5	2	20
	WSNO	18	0	10	10	18
	SAP	21	0	0	0	21
	Total	63	0	15	12	59
MIPLIB	Total	252	129	184	193	140
PB16	Factor	192	186	186	186	182
	Kexu	40	14	17	14	40
	Logic synthesis	74	72	72	72	74
	Market split	40	16	22	12	8
	Mps	35	30	33	33	25
	Numerical	34	13	19	25	8
	Prime	156	127	130	131	132
	Reduced mps	273	100	150	160	43
	Total	1524	995	1198	1201	1107

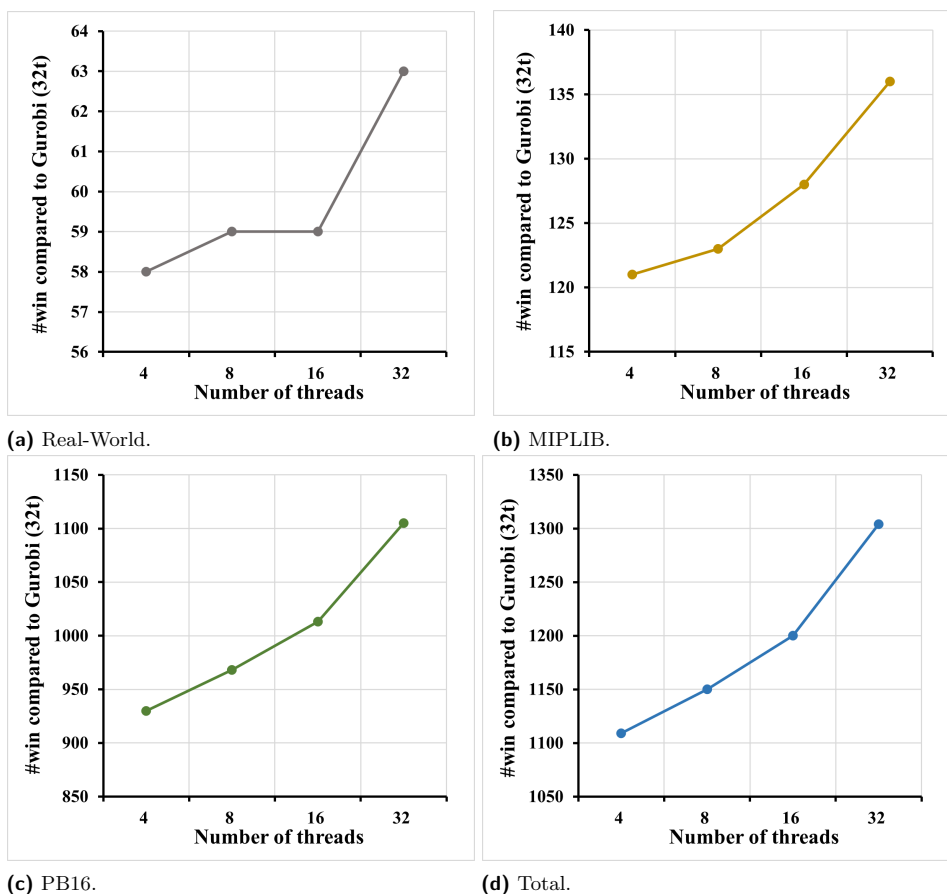
6 Conclusions

We proposed two local search solvers for the PBO problem: *DLS-PBO* and *ParLS-PBO*. *DLS-PBO* is an enhanced version of the *LS-PBO* solver, incorporating a dynamic scoring mechanism. *ParLS-PBO* is a parallel solver with a solution pool collecting good solutions from multiple threads. The solution pool guides the local search process by providing better starting points and utilizing polarity information from high-quality solutions to improve the scoring function. Experimental results show that our parallel solver has significantly better performance than all sequential solvers and exhibits strong competitiveness against the parallel versions of *Gurobi*.

The ideas of this work can be applied to other problems, particularly including SAT and MaxSAT. It is also interesting to implement a distributed version of *ParLS-PBO* for cloud computation.

■ **Table 5** Performance evaluation between *ParLS-PBO* and its variants.

Benchmark	#Ins	V_1 vs. <i>ParLS-PBO</i>				V_2 vs. <i>ParLS-PBO</i>			
		V_1		<i>ParLS-PBO</i>		V_2		<i>ParLS-PBO</i>	
		#win	avg _{sc} *	#win	avg _{sc} *	#win	avg _{sc} *	#win	avg _{sc} *
cutoff=300s									
Real-World	63	28	0.993	52	0.998	36	0.996	44	0.998
MIPLIB	252	137	0.857	199	0.868	166	0.871	185	0.870
PB16	1524	1095	0.835	1231	0.847	1168	0.841	1182	0.844
cutoff=3600s									
Real-World	63	19	0.981	62	1.0	37	0.995	48	0.999
MIPLIB	252	145	0.870	203	0.897	181	0.897	180	0.893
PB16	1524	1111	0.840	1251	0.854	1191	0.852	1214	0.852



■ **Figure 2** Scalability Analysis (Time limit is set to 300).

References

- 1 Adrian Balint and Uwe Schöning. Choosing probability distributions for stochastic local search and the role of make versus break. In Alessandro Cimatti and Roberto Sebastiani, editors, *Theory and Applications of Satisfiability Testing - SAT 2012 - 15th International Conference, Trento, Italy, June 17-20, 2012. Proceedings*, volume 7317 of *Lecture Notes in Computer Science*, pages 16–29. Springer, 2012. doi:10.1007/978-3-642-31612-8_3.
- 2 Peter Barth. A davis-putnam based enumeration algorithm for linear pseudo-boolean optimization. Technical report, Max Plank Institute for Computer Science, 1995.
- 3 Belaid Benhamou, Lakhdar Sais, and Pierre Siegel. Two proof procedures for a cardinality based language in propositional calculus. In Patrice Enjalbert, Ernst W. Mayr, and Klaus W. Wagner, editors, *STACS 94, 11th Annual Symposium on Theoretical Aspects of Computer Science, Caen, France, February 24-26, 1994, Proceedings*, volume 775 of *Lecture Notes in Computer Science*, pages 71–82. Springer, 1994. doi:10.1007/3-540-57785-8_132.
- 4 Jeremias Berg, Emilia Oikarinen, Matti Järvisalo, and Kai Puolamäki. Minimum-width confidence bands via constraint optimization. In *International Conference on Principles and Practice of Constraint Programming*, pages 443–459. Springer, 2017.
- 5 Ksenia Bestuzheva, Mathieu Besançon, Wei-Kun Chen, Antonia Chmiela, Tim Donkiewicz, Jasper van Doornmalen, Leon Eifler, Oliver Gaul, Gerald Gamrath, Ambros Gleixner, et al. The scip optimization suite 8.0. *arXiv preprint arXiv:2112.08872*, 2021.
- 6 Shaowei Cai, Chuan Luo, and Kaile Su. Ccanr: A configuration checking based local search solver for non-random satisfiability. In Marijn Heule and Sean A. Weaver, editors, *Theory and Applications of Satisfiability Testing - SAT 2015 - 18th International Conference, Austin, TX, USA, September 24-27, 2015, Proceedings*, volume 9340 of *Lecture Notes in Computer Science*, pages 1–8. Springer, 2015. doi:10.1007/978-3-319-24318-4_1.
- 7 Yuning Chen and Jin-Kao Hao. Memetic search for the generalized quadratic multiple knapsack problem. *IEEE Trans. Evol. Comput.*, 20(6):908–923, 2016. doi:10.1109/TEVC.2016.2546340.
- 8 Zhihan Chen, Xindi Zhang, Yuhang Qian, and Shaowei Cai. Prs: A new parallel/distributed framework for sat. *SAT COMPETITION 2023*, page 39, 2023.
- 9 Yi Chu, Shaowei Cai, and Chuan Luo. Nuwls: Improving local search for (weighted) partial maxsat by new weighting techniques. In Brian Williams, Yiling Chen, and Jennifer Neville, editors, *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, pages 3915–3923. AAAI Press, 2023. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/25505>, doi:10.1609/AAAI.V37I4.25505.
- 10 Yi Chu, Shaowei Cai, Chuan Luo, Zhendong Lei, and Cong Peng. Towards more efficient local search for pseudo-boolean optimization. In *29th International Conference on Principles and Practice of Constraint Programming (CP 2023)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2023.
- 11 Olivier Coudert and Jean Christophe Madre. New ideas for solving covering problems. In Bryan Preas, editor, *Proceedings of the 32st Conference on Design Automation, San Francisco, California, USA, Moscone Center, June 12-16, 1995*, pages 641–646. ACM Press, 1995. doi:10.1145/217474.217603.
- 12 Jo Devriendt, Stephan Gocht, Emir Demirovic, Jakob Nordström, and Peter J. Stuckey. Cutting to the core of pseudo-boolean optimization: Combining core-guided search with cutting planes reasoning. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 3750–3758. AAAI Press, 2021. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/16492>, doi:10.1609/AAAI.V35I5.16492.
- 13 Niklas Eén and Niklas Sörensson. Translating pseudo-boolean constraints into sat. *Journal on Satisfiability, Boolean Modeling and Computation*, 2(1-4):1–26, 2006.

- 14 Jan Elffers and Jakob Nordström. Divide and conquer: Towards faster pseudo-boolean solving. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 1291–1299. ijcai.org, 2018.
- 15 Armin Biere Katalin Fazekas Mathias Fleury and Maximilian Heisinger. Cadical, kissat, paracooba, plingeling and treengeling entering the sat competition 2020. *SAT COMPETITION*, 2020:50, 2020.
- 16 Ludovic Le Frioux, Souheib Baarir, Julien Sopena, and Fabrice Kordon. Painless: A framework for parallel SAT solving. In Serge Gaspers and Toby Walsh, editors, *Theory and Applications of Satisfiability Testing - SAT 2017 - 20th International Conference, Melbourne, VIC, Australia, August 28 - September 1, 2017, Proceedings*, volume 10491 of *Lecture Notes in Computer Science*, pages 233–250. Springer, 2017. doi:10.1007/978-3-319-66263-3_15.
- 17 LLC Gurobi Optimization. Gurobi optimizer reference manual, 2021.
- 18 Markus Iser, Jeremias Berg, and Matti Järvisalo. Oracle-based local search for pseudo-boolean optimization. In Kobi Gal, Ann Nowé, Grzegorz J. Nalepa, Roy Fairstein, and Roxana Radulescu, editors, *ECAI 2023 - 26th European Conference on Artificial Intelligence, September 30 - October 4, 2023, Kraków, Poland - Including 12th Conference on Prestigious Applications of Intelligent Systems (PAIS 2023)*, volume 372 of *Frontiers in Artificial Intelligence and Applications*, pages 1124–1131. IOS Press, 2023. doi:10.3233/FAIA230387.
- 19 Luyu Jiang, Dantong Ouyang, Qi Zhang, and Liming Zhang. Decils-pbo: an effective local search method for pseudo-boolean optimization. *CoRR*, abs/2301.12251, 2023. doi:10.48550/arXiv.2301.12251.
- 20 Gergely Kovásznai, Balázs Erdélyi, and Csaba Biró. Investigations of graph properties in terms of wireless sensor network optimization. In *2018 IEEE International Conference on Future IoT Technologies (Future IoT)*, pages 1–8. IEEE, 2018.
- 21 Gergely Kovásznai, Krisztián Gajdár, and Laura Kovács. Portfolio sat and smt solving of cardinality constraints in sensor network optimization. In *2019 21st International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC)*, pages 85–91. IEEE, 2019.
- 22 Daniel Le Berre and Anne Parrain. The sat4j library, release 2.2. *Journal on Satisfiability, Boolean Modeling and Computation*, 7(2-3):59–64, 2010.
- 23 Zhendong Lei and Shaowei Cai. Solving (weighted) partial maxsat by dynamic local search for SAT. In Jérôme Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 1346–1352. ijcai.org, 2018. doi:10.24963/ijcai.2018/187.
- 24 Zhendong Lei, Shaowei Cai, Chuan Luo, and Holger H. Hoos. Efficient local search for pseudo boolean optimization. In Chu-Min Li and Felip Manyà, editors, *Theory and Applications of Satisfiability Testing - SAT 2021 - 24th International Conference, Barcelona, Spain, July 5-9, 2021, Proceedings*, volume 12831 of *Lecture Notes in Computer Science*, pages 332–348. Springer, 2021. doi:10.1007/978-3-030-80223-3_23.
- 25 Chu Min Li and Yu Li. Satisfying versus falsifying in local search for satisfiability - (poster presentation). In Alessandro Cimatti and Roberto Sebastiani, editors, *Theory and Applications of Satisfiability Testing - SAT 2012 - 15th International Conference, Trento, Italy, June 17-20, 2012. Proceedings*, volume 7317 of *Lecture Notes in Computer Science*, pages 477–478. Springer, 2012. doi:10.1007/978-3-642-31612-8_43.
- 26 Stan Y. Liao and Srinivas Devadas. Solving covering problems using lpr-based lower bounds. In Ellen J. Yoffa, Giovanni De Micheli, and Jan M. Rabaey, editors, *Proceedings of the 34th Conference on Design Automation, Anaheim, California, USA, Anaheim Convention Center, June 9-13, 1997*, pages 117–120. ACM Press, 1997. doi:10.1145/266021.266046.
- 27 Marius Lindauer, Katharina Eggenesperger, Matthias Feurer, André Biedenkapp, Difan Deng, Carolin Benjamins, Tim Ruhkopf, René Sass, and Frank Hutter. Smac3: A versatile bayesian optimization package for hyperparameter optimization. *Journal of Machine Learning Research*, 23(54):1–9, 2022.

- 28 Joao Marques-Silva, Inês Lynce, and Sharad Malik. Conflict-driven clause learning sat solvers. In *Handbook of satisfiability*, pages 133–182. IOS Press, 2021.
- 29 Ruben Martins, Vasco Manquinho, and Inês Lynce. Exploiting cardinality encodings in parallel maximum satisfiability. In *2011 IEEE 23rd International Conference on Tools with Artificial Intelligence*, pages 313–320. IEEE, 2011.
- 30 Ruben Martins, Vasco Manquinho, and Inês Lynce. Parallel search for maximum satisfiability. *AI Communications*, 25(2):75–95, 2012.
- 31 Ruben Martins, Vasco M. Manquinho, and Inês Lynce. Open-wbo: A modular maxsat solver,. In Carsten Sinz and Uwe Egly, editors, *Theory and Applications of Satisfiability Testing - SAT 2014 - 17th International Conference, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 14-17, 2014. Proceedings*, volume 8561 of *Lecture Notes in Computer Science*, pages 438–445. Springer, 2014. doi:10.1007/978-3-319-09284-3_33.
- 32 Ruben Martins and Justine Sherry. Lisbon wedding: seating arrangements using maxsat. *MaxSAT Evaluation*, pages 25–26, 2017.
- 33 Olivier Roussel and Vasco M. Manquinho. Pseudo-boolean and cardinality constraints. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, pages 1087–1129. IOS Press, 2021. doi:10.3233/FAIA201012.
- 34 Yuji Shinano, Stefan Heinz, Stefan Vigerske, and Michael Winkler. Fiberscip - A shared memory parallelization of SCIP. *INFORMS J. Comput.*, 30(1):11–30, 2018. doi:10.1287/ijoc.2017.0762.
- 35 Pavel Smirnov, Jeremias Berg, and Matti Järvisalo. Improvements to the implicit hitting set approach to pseudo-boolean optimization. In Kuldeep S. Meel and Ofer Strichman, editors, *25th International Conference on Theory and Applications of Satisfiability Testing, SAT 2022, August 2-5, 2022, Haifa, Israel*, volume 236 of *LIPICs*, pages 13:1–13:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICs.SAT.2022.13.
- 36 Rodrigue Konan Tchinda and Clémentin Tayou Djamegni. Hkis, head, pakis and painless exmaplelcmdistchronobt in the sc21. *SAT COMPETITION*, 2021:26, 2021.