Improving Watched Pseudo-Boolean Propagation with Significant Literals

- ₃ Mia Müßig ⊠©
- 4 Institut für Informatik, Ludwig-Maximilians-Universität München, Germany
- 5 Department of Statistical Learning, ScaDS.AI Leipzig, Germany
- 🛛 Jan Johannsen 🖂
- 7 Institut für Informatik, Ludwig-Maximilians-Universität München, Germany

8 — Abstract -

A key challenge in pseudo-boolean solving is the efficient detection and processing of unit literals. 9 In SAT solving this is done by using the watched literal scheme, but for general pseudo-boolean 10 constraints there is no dominant method. This paper introduces the significant literal framework 11 to generalize existing watched literal schemes for pseudo-boolean solvers, which is implemented 12 in a modification of the ROUNDINGSAT solver. For this modification, small improvements can be 13 observed on the instances from the decision and optimization tracks of the 2024 Pseudo-Boolean 14 Competition, and a substantial improvement in instances with large coefficients like Knapsack. 15 2012 ACM Subject Classification Theory of computation 16

- 17 Keywords and phrases Pseudo-Boolean Solving, Unit Propagation, Watched Literals
- ¹⁸ Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

Acknowledgements This paper is based on the first author's Master's Thesis [8] supervised by the second author.

21 Introduction

Contemporary satisfiability (SAT) solvers achieve efficiency by conflict driven clause learning 22 (CDCL) and unit propagation, making them a powerful tool for boolean problems like 23 model checking in chip development, formal software verification and scheduling problems. 24 On the other hand, Integer Linear Programming (ILP) solvers use integer relaxation and 25 the branch-and-cut procedure to dominate in industrial applications including production 26 planning, network optimization and portfolio selection. Pseudo-boolean solving has emerged 27 as a promising middle ground between the two approaches, aiming to combine the advantages 28 of both methods. SAT methods like unit propagation and conflict analysis are now applied 29 to the powerful cutting plane system underlying ILP solvers. 30

However, neither unit propagation nor conflict analysis can be generalized to arbitrary 31 pseudo-boolean constraints without careful considerations. The latter received significant 32 attention, with various papers discussing different methods of constructing conflict constraints 33 [2, 5, 13]. For unit propagation, the focus is on adapting the watched literal scheme introduced 34 in the SAT solver Chaff [7], which is part of almost all modern CDCL solvers. However, 35 many teams developing competitive pseudo-boolean solvers observed none to only minimal 36 improvement compared to the simpler counting method [1, 2, 13]. This changed when 37 Devriendt [3] obtained a significant performance increase for his watched literal scheme, 38 which replaces the computation of the maximum coefficient of unassigned literals used to 39 determine the watched literal set by a fixed upper bound. This method has been implemented 40 in the ROUNDINGSAT solver, which is currently considered the fastest pseudo-boolean solver 41 [4]. Recently it has been observed that the performance improvement of watched literals can 42 be amplified by caching optimizations and a hybrid unit propagation approach [10]. 43

© Mia Müßig and Jan Johannsen; licensed under Creative Commons License CC-BY 4.0 42nd Conference on Very Important Topics (CVIT 2016). Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:13 Leibniz International Proceedings in Informatics LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

23:2 Improving Watched Pseudo-Boolean Propagation with Significant Literals

This paper aims to demonstrate that the watched literal scheme of ROUNDINGSAT can 44 be further improved by introducing significant literals, which are especially effective for 45 constraints with large coefficient sizes. Our approach uses a tighter upper bound by finding 46 the maximum coefficient of unassigned *significant* literals, but keeps the fixed upper bound as 47 in Devriendt's scheme [3] for constraints without significant literals. The significance criterion 48 for literals is determined in such a way that we strike a balance between time savings from 49 having fewer watched literals and the higher cost of updating the new watched literal scheme. 50 Additionally, we attempt to explain a connection of the performance differences between the 51 various watched literal schemes and the size distributions of the constraint coefficients. 52

⁵³ 2 Preliminaries

A pseudo-boolean problem consists of variables $x_i \in \{0, 1\}$, with literals l_i representing either x_i or $\overline{x}_i := 1 - x_i$ and pseudo-boolean constraints C of the form $\sum_{i=1}^n a_i l_i \ge b$ for $a_i, b \in \mathbb{N}$. Without loss of generality we will assume that the coefficients are in descending order, so $\forall i : a_i \ge a_{i+1}$. A cardinality constraint is a pseudo-boolean constraint with $\forall i : a_i = 1$. If additionally b = 1, the constraint is a clause.

We denote the current (partial) assignment of the variables l_i with ρ , identified with the set of literals set to 1. The *slack* of a pseudo-boolean constraint with the current assignment ρ is defined as $slack(C, \rho) = -b + \sum_{\bar{l}_i \notin \rho} a_i$.

Informally, the slack represents the amount by which the left-hand side can exceed the right-hand side without unassigning literals. Thus, if $slack(C, \rho) < 0$, the constraint C can not be satisfied with ρ , and we need to backtrack.

⁶⁵ A literal l_i is called a *unit literal*, if $slack(C, \rho) < a_i$. This means that l_i must be set to 1 ⁶⁶ in order to satisfy the constraint C, as otherwise $slack(C, \rho \cup \{\overline{l_i}\}) < 0$.

For a pseudo-boolean constraint C we denote its watched literals with W(C), a subset of its non-falsified literals, i.e. $W(C) \subseteq \{l_i \in C : \overline{l_i} \notin \rho\}$. This is different from SAT, where only unassigned literals are watched. The maximum coefficient a_{max} of a constraint C is defined as $a_{max} := \max\{a_i \in C : l_i, \overline{l_i} \notin \rho\}$. If all literals of C are assigned, we define it as $a_{max} = 0$. The watch slack is the slack restricted to W(C): $wslack(C, \rho) = -b + \sum_{l_i \in W(C)} a_i$. By definition $wslack(C, \rho) \leq slack(C, \rho)$ holds.

⁷³ ► Lemma 1. C contains a unit literal if and only if no watched literals set W(C) exists such ⁷⁴ that $wslack(C, ρ) \ge a_{max}$.

⁷⁵ **Proof.** If C contains a unit literal, then by definition we have an unassigned literal l_i with ⁷⁶ $slack(C, \rho) < a_i$ and $a_{max} \ge a_i$. So we arrive at $a_{max} \ge a_i > slack(C, \rho) \ge wslack(C, \rho)$. ⁷⁷ Thus, no watched literal set W(C) with $wslack(C, \rho) \ge a_{max}$ can exist.

If for all sets of literals W(C) we have $wslack(C, \rho) < a_{max}$, this also holds true for $W(C) = \{l_i | \overline{l_i} \notin \rho\}$, i.e. if we watch all non-falsified literals. In that case we have $wslack(C, \rho) = slack(C, \rho)$, leading us to $slack(C, \rho) < a_{max}$, and therefore the literal corresponding to a_{max} must be unit.

If we replace a_{max} with an upper bound, a weaker version of Lemma 1 still holds:

Lemma 2. Let a_{up} be an upper bound, i.e. $a_{up} \ge a_{max}$. If C contains a unit literal, then no watched literals set W(C) exists such that $wslack(C, \rho) \ge a_{up}$.

Proof. By Lemma 1 we know that if C contains a unit literal, then we can not choose watched literals W(C) such that $wslack(C, \rho) \ge a_{max}$. Since $a_{up} \ge a_{max}$, this immediately implies that we also can not find W(C) with $wslack(C, \rho) \ge a_{up}$.

The problem of using Lemma 1 as a criterion for finding unit literals is that a_{max} needs to be known for all constraints. So every time a variable x_i is assigned or unassigned, we have to update a_{max} in all constraints that contain x_i or \overline{x}_i . This procedure can take up to two thirds of the total unit propagation runtime [13], and defeats the original purpose of watched literals.

ROUNDINGSAT instead uses a criterion based on Lemma 2 with $a_{up} = a_1$ [3]. As a_1 is constant, no updates are necessary to preserve $a_{up} \ge a_{max}$. One downside of this approach is that now our unit propagation methods can encounter false positives, as the converse of Lemma 2 does not hold true. Additionally, it often leads to a larger W(C), increasing the work necessary for maintaining the watched literal scheme.

3 Watched Literal Propagation with Significant Literals

⁹⁹ The idea behind significant literals is finding a middle ground between always updating a_{max} ¹⁰⁰ and replacing it with a constant upper bound. For that we choose an arbitrary criterion ¹⁰¹ determining if l_i is significant for a constraint C. We denote the set of all significant literals ¹⁰² for C with S(C). Since the criterion does not depend on ρ , S(C) is constant.

Now we define $a_{smax} := \max\{a_i \in C : l_i, \overline{l_i} \notin \rho \lor l_i \notin S(C)\}$, i.e. a_{smax} is the largest coefficient of a literal which is either unassigned or not significant. As with a_{max} , if all literals are assigned and significant, we simply define $a_{smax} = 0$. By definition $a_{smax} \ge a_{max}$, so we can apply Lemma 2 with $a_{up} = a_{smax}$.

This framework generalizes the two previously mentioned methods. If we choose a significance criterion which declares all literals as significant, we have $a_{smax} = a_{max}$ and recover the method which calculates a_{max} for all constraints. If we instead declare no literal to be significant for any constraint, we have $a_{smax} = a_1$ and recover the current method in ROUNDINGSAT [3].

112 3.1 Algorithm

Algorithm 1 demonstrates how to set $C.a_{smax}$ to the value of a_{smax} . Here C.l[k] is a reference to a literal l_k occurring in C, C.a[k] denotes its corresponding coefficient, and size(C) denotes the number of literals in C. In the while loop we search for an unassigned or non-significant literal, and finish immediately if we found one. Since the literals in C are sorted by descending order of their coefficients, the first found literal is guaranteed to have the largest coefficient, and we have found a_{smax} .

Algorithm 1 SETSMAX (constraint C)

```
1: Input: Constraint C
 2: if C.lastset < bkjmps then
         C.k \leftarrow 1
 3:
         C.lastset \leftarrow bkjmps
 4:
 5: while C.k \leq \text{size}(C) do
         l_k \leftarrow C.l[k]
 6:
         if l_k, \bar{l}_k \notin \rho or l_i \notin S(C) then
 7:
              C.a_{smax} \leftarrow C.a[k]
 8:
 9:
              return
         C.k \leftarrow C.k + 1
10:
11: C.a_{smax} \leftarrow 0
```



23:4 Improving Watched Pseudo-Boolean Propagation with Significant Literals

To further optimize this procedure, we remember the index k at which we left the loop in the last call to SETSMAX. If no backtracking occurred in the meantime, no literals have been unassigned and so all literals before k are still assigned and significant. Thus, the first condition in algorithm 1 ensures that we only restart the search from the beginning if a backjump occurred. This idea is directly inspired by the algorithm 2 first presented by Devriendt [3].

¹²⁵ When we now assign a literal we check in each constraint in which it is significant, if its ¹²⁶ corresponding coefficient is equal to a_{smax} of that constraint. Only if that it is the case, we ¹²⁷ need to call algorithm 1 to update the value. For unassigning we only need to check for each ¹²⁸ significant constraint if its corresponding coefficient is bigger than a_{smax} and if so set a_{smax} ¹²⁹ to its coefficient.

¹³⁰ Now we need to integrate the new $C.a_{smax}$ into the unit propagation of ROUNDINGSAT, ¹³¹ which consists of PROPAGATEOPT, PROCESSWATCHES and BACKJUMP [3]. Luckily only the ¹³² first routine needs to be modified into algorithm 2, since it is the only one dependent on the ¹³³ choice of a_{up} . This independence of a_{up} also holds for proof of the watch slack invariant [3], ¹³⁴ which means that all routines correctly preserve $C.wslk = wslack(C, \rho)$.

Algorithm 2 PROPAGATE (constraint C, integer idx), modification of PROPAGATEOPT [3]

1: Input: assignment ρ , literal l, constraint C 2: if C.lastprop < bkjmps then $C.i \leftarrow 1$ 3: $C.j \leftarrow 1$ 4: C.lastprop $\leftarrow bkjmps$ 5:if $C.wslk + C.a[idx] \ge C.a_{smax} \lor S(C) \neq \emptyset$ then 6: while $C.i \leq \text{size}(C)$ and $C.\text{wslk} < C.a_{smax}$ do 7: if $\overline{C.l}[i] \notin \rho$ and $C.l[i] \notin W(C)$ then 8: \triangleright Add C.l[i] to watched literals 9: $W(C) \leftarrow W(C) \cup \{C.l[i]\}$ 10: $C.wslk \leftarrow C.wslk + C.a[i]$ 11: $C.i \leftarrow C.i + 1$ 12:13: if C.wslk $\geq C.a_{smax}$ then ▷ Enough watched literals, unwatch propagated literal 14: $W(C) \leftarrow W(C) \setminus \{C.l[idx]\}$ 15:return OK 16:17: if C.wslk < 0 then return CONFLICT 18: \triangleright A unit literal could exist 19: while $C.j \leq \text{size}(C)$ and C.wslk < C.a[j] do 20: $l_i \leftarrow C.l[j]$ 21: if $l_i, \bar{l}_i \notin \rho$ then \triangleright Enqueue new unit literal 22: 23: $\rho \leftarrow \rho \cup \{l_j\}$ 24: $C.j \leftarrow C.j + 1$ 25: return OK

However, one needs to be careful about which optimization of PROPAGATEOPT are still logically sound for a non-static a_{up} . The first optimization can be kept, as the following lemma shows.

138 **Lemma 3.** If no backtracking occurred, the while loops in line 7 and 19 can be restarted at

¹³⁹ the index where they left off in the last call, without changing the behaviour of the algorithm.

¹⁴⁰ **Proof.** Without backtracking any literal assigned in ρ at the last call to function 2 remain ¹⁴¹ assigned. Thus, if *i* is the index, where the while loop of line 7 left off, the literals $l_1, \ldots l_{i-1}$ ¹⁴² either remain false or are already watched, and we can safely skip them.

Similarly, if the while loop of line 19 terminates with index j, then all literals $l_1, \ldots l_{j-1}$ are already assigned. Without backtracking this still holds true now, and we can restart the search with index j.

The second optimization, which skips the search for new watched literals if C.wslk + 146 $C.a[idx] \geq C.a_{up}$, does not work for non-constant a_{up} . The idea behind it is that if at some 147 assignment the while loop of line 7 terminates because of the condition C.wslk $< C.a_{smax}$, 148 all non-falsified literals are watched. Until we have backtracked so far that this specific 149 assignment is reversed, all non-falsified literals remain watched, and any further search 150 for new literals to watch will be unnecessary. However, with a non-constant a_{up} further 151 assignments can make it possible to fulfil $wslack(C, \rho) \geq a_{up}$, even if that was not possible at 152 some past call. Thus, the optimization can only be applied if a_{up} is constant in the constraint. 153 In our case we simply check this when $S(C) = \emptyset$, since we have $a_{smax} = a_1$ if no literals are 154 significant for a constraint. 155

3.2 Significance Criteria

¹⁵⁷ To elaborate what significance criteria are useful, we look at the following "worst case" ¹⁵⁸ constraint for the current ROUNDINGSAT unit propagation system: $100y + \sum_{i=1}^{100} x_i \ge 10$.

Starting with $\rho = \emptyset$, we have $a_1 = a_{max} = 100$ and $slack(C, \rho) = 190$. Then the implementation iteratively adds watched literals until $wslack(C, \rho) \ge 100$, which results in $W(C) = \{y, x_1, \ldots x_{10}\}$. If the next decision of the solver is $\rho = \{\overline{y}\}$, all 100 x_i literals will be added to W(C) without achieving $wslack(C, \rho) \ge 100$. This means that until the solver reverses the assignment of y, the watched set W(C) will be unnecessarily large, and no unit literal can exist until at least 90 of the x_i variables are assigned.

When we instead use significant literals and a criterion which declares y to be significant for the constraint, we only need $wslack(C, \rho) \ge 1$ after $\rho = \{\overline{y}\}$. This allows for the much smaller watched set $W(C) = \{x_1, \ldots, x_{11}\}$, which reduces the workload for updating the watched literal scheme in further assignments.

Here the difference between the two watched literal schemes is exaggerated, as the example constraint is deliberately chosen to amplify this problem. Simply choosing the equivalent constraint $10y + \sum_{i=1}^{100} x_i \ge 10$ allows the current ROUNDINGSAT implementation to watch only roughly twice as many literals as the significant literal approach. However, it still illustrates that the difference between the two watched literal schemes is mainly dependent on the size of the constraint coefficients. This will also be supported by empirical evidence in the following section.

We experiment with various significance criteria, which aim to identify literals where the increased cost of updating a_{smax} is outweighed by allowing for a smaller watched literal set W(C). Let $c \in \mathbb{N}$ be a fixed cut-off value, and $s \in \mathbb{N}$ a fixed scaling factor. We consider the following criteria, where the literal l_i with coefficient a_i is significant for a constraint C if: (Absolute size) $a_i > c$

- (Absolute size of maximum coefficient) $a_1 > c$
- (Relative size) $a_1 > s \sum_{i=2}^n a_i$

Each of the three criteria can also be restricted to only input constraints, which we will mark as "C input" in the legends of runtime plots.



Figure 1 Runtime comparison on the DEC-LIN track. The number of instances solved is given in parentheses.

4 Experimental Evaluation

To perform the evaluation we use commit d34b6bed of the ROUNDINGSAT solver, which 186 is the latest version as of December 2024 [4]. We extended ROUNDINGSAT to implement 187 the unit propagation with significant literals. The implementation and the data presented 188 in the following plots are publicly available [9]. It should be noted that - as in the current 189 ROUNDINGSAT implementation - the respective watched literal scheme is only applied 190 to "true" pseudo-boolean constraints, while clauses and cardinality constraints are treated 191 separately. All runtime measurements are done without proof logging, although we verified 192 all certificates of an independent run using the VERIPB verifier [6]. The benchmark was run 193 on an AMD Ryzen 5950X CPU with a timeout of 3600s. 194

Our first dataset will consist of the 398 selected instances of the DEC-LIN track in 195 the Pseudo-Boolean Competition 2024 [11], which aim to be a representative sample of 196 pseudo-boolean decision problems. We evaluate the unmodified ROUNDINGSAT solver (R-197 SAT) without the optional SOPLEX Linear Programming integration and without the hybrid 198 mode suggested by Robert Nieuwenhuis et al. [10], as both actually reduce the number of 199 instances solved for this specific dataset. Additionally, we evaluate the counting method -200 which explicitly calculates the slack of each constraint - by disabling watched propagation in 201 ROUNDINGSAT. Finally we present the best performing significance criteria and "Standard 202 Watched", which uses a significance criterion declaring every literal to be significant for all 203 its constraints, in order to represent the traditional watched literal scheme with $a_{up} = a_{max}$. 204 Figure 1 demonstrates that the counting method is less efficient than the unit propagation 205 by Devriendt [3], although faster than the naive way of implementing watched literals. We 206 also observe that some significance criteria already gain a small advantage over the existing 207 ROUNDINGSAT implementation on the DEC-LIN instances. An extensive evaluation of 208 different significance criteria on these instances is given in appendix C. 209





Figure 2 Distribution of coefficients in input constraints.

Figure 3 Distribution of coefficients in learned constraints.

As shown in appendix A, the most successful criteria for DEC-LIN also perform well on the OPT-LIN track in the Pseudo Boolean Competition 2024, which comprises a collection of 487 optimization problems. Here again, a small improvement compared to ROUNDINGSAT can be observed.

Across both datasets we notice that unit propagation with significant literals works best when it is only applied to the input constraints and not to the learned constraints. Our explanation for this behaviour is the difference in the distribution of coefficient sizes as shown in Figures 2 and 3.

To produce Figure 2 we have collected all constraints in the selected instances of the DEC-LIN track, which are not clauses or cardinality constraints and group all their coefficients by absolute values in buckets. For Figure 3 we did the same for all learned constraints ROUNDINGSAT on the dataset, but excluded instances which are not solved within the 3600s timeout.

We observe that coefficients in input constraints are far more unevenly distributed. Thus, a_{smax} can often be far smaller than a_{max} , leading to a smaller and therefore less expensive watched literal set. On learned constraints with a more even distribution, a_{smax} remains similar to a_{max} , thus the cost of updating a_{smax} outweighs the benefit of only marginally smaller watched literal sets.

While significant literals already yield a small improvement on general instances, they are substantially better if the instances mainly consist of constraints with large coefficients. One example are the 783 Knapsack instances submitted to the OPT-LIN track of the Pseudo Boolean Competition 2024, as Figure 4 shows.

The plot demonstrates that the two criteria clearly outperform the unmodified ROUND-INGSAT solver, both in the number of instances solved and the general runtime of hard instances.

In Appendix B we further support this observation by looking at another application with large coefficient in the benchmark dataset of the Pseudo Boolean Competition, the "Virtual Machine Consolidation" problem [12]. As Figure 6 shows, the significane criteria, which performed best for the Knapsack instances, also outperform the current ROUNDINGSAT implementation there.

23:8 Improving Watched Pseudo-Boolean Propagation with Significant Literals



Figure 4 Runtime comparison of significant literal schemes on Knapsack instances. The number of instances solved is given in parentheses.

240 **5** Conclusion

It has been successfully demonstrated how significant literals can speedup unit propagation
for pseudo-boolean constraints. Especially for applications with a large average coefficient
size we have observed a substantial improvement compared to the current ROUNDINGSAT
watched propagation scheme.

In addition, an avenue for future improvements is indicated by the observation that the performance differences between the watched literal schemes is mainly determined by the distribution of coefficients. One idea is to choose a conflict analysis method specifically based on its suitability for the watched literal scheme. Alternatively, the significance criterion itself could be selected during the preprocessing step based on the coefficients of the individual instance.

²⁵¹ — References

- Fadi A. Aloul, Arathi Ramani, Igor L. Markov, and Karem A. Sakallah. Generic ilp versus specialized 0-1 ilp: an update. In *Proceedings of the 2002 IEEE/ACM International Conference* on Computer-Aided Design (ICCAD '02), pages 450–457, New York, NY, USA, 2002. ACM. doi:10.1145/774572.774638.
- Donald Chai and Andreas Kuehlmann. A fast pseudo-boolean constraint solver. In *Proceedings* of the 40th Annual Design Automation Conference, DAC '03, page 830–835, 2003. doi: 10.1145/775832.776041.
- Jo Devriendt. Watched Propagation of 0-1 Integer Linear Constraints. In Principles and Practice of Constraint Programming, pages 160–176, 2020.
- Jan Elffers, Jo Devriendt, Stephan Gocht, and Jakob Nordström. RoundingSat the pseudo boolean solver powered by proof complexity!, 2024. URL: https://gitlab.com/MIAOresearch/
 software/roundingsat.

- Jan Elffers and Jakob Nordström. Divide and conquer: Towards faster pseudo-boolean solving.
 In International Joint Conference on Artificial Intelligence, 2018.
- 6 Stephan Gocht, Ciaran McCreesh, and Jakob Nordström. VeriPB: the easy way to make your combinatorial search algorithm trustworthy. In From Constraint Programming to Trustworthy AI, workshop at the 26th International Conference on Principles and Practice of Constraint Programming (CP '20), September 2020.
- 7 Matthew W. Moskewicz, Lintao Zhang, Conor F. Madigan, Sharad Malik, and Ying Zhao.
 Chaff: Engineering an Efficient SAT Solver . In *Design Automation Conference*, pages 530–535, 2001. doi:10.1109/DAC.2001.935565.
- 8 Mia Müßig. Investigating watched literal schemes for pseudo-boolean solvers. Master's thesis,
 LMU Munich, 2025.
- 9 Mia Müßig. RoundingSAT with significant literals, 2025. URL: https://gitlab2.cip.ifi.
 1mu.de/muessig/roundingsat.
- Robert Nieuwenhuis, Albert Oliveras, Enric Rodríguez-Carbonell, and Rui Zhao. Speeding up
 Pseudo-Boolean Propagation. In 27th International Conference on Theory and Applications
 of Satisfiability Testing (SAT 2024), volume 305 of Leibniz International Proceedings in
 Informatics (LIPIcs), pages 22:1–22:18, 2024. doi:10.4230/LIPIcs.SAT.2024.22.
- Pseudo-Boolean Competition. Pseudo-Boolean Competition 2024, 2024. Organized by Olivier
 Roussel, CRIL, Université d'Artois, France. URL: https://www.cril.univ-artois.fr/PB24/.
- Bruno César Ribas, Rubens Massayuki Suguimoto, Razer A.N.R. Montaño, Fabiano Silva, and Marcos Castilho. PBFVMC: A new pseudo-boolean formulation to virtual-machine consolidation. In 2013 Brazilian Conference on Intelligent Systems, pages 201–206, 2013. doi:10.1109/BRACIS.2013.41.
- Hossein M. Sheini and Karem A. Sakallah. Pueblo: a modern pseudo-boolean sat solver. In
 Design, Automation and Test in Europe, volume 2, pages 684–685, 2005. doi:10.1109/DATE.
 2005.246.
- ²⁹⁰ A Comparison with OPT-LIN Instances

The comparison for the OPT-LIN track of the Pseudo-Boolean Competition 2024 is shown in figure 5.

²⁹³ **B** Comparison with PBFVMC Instances

Figure 6 shows the runtime comparison of significant literal schemes on the PBFVMC instances. The dataset includes 27 decision and 27 optimization problems submitted as benchmarks for the Pseudo Boolean Competition 2015. Since these are far fewer instances than for Knapsack, it makes it harder to draw decisive conclusion from this data alone.

298 **C**

C Comparison of different Significance Definitions

In figures 7 through 12, we compare the runtime of different significance criteria on the
 DEC-LIN instances of the Pseudo-Boolean Competition 2024.



Figure 5 Runtime comparison of significant literal schemes on the OPT-LIN Track.



Figure 6 Runtime comparison of significant literal schemes on the PBFVMC instances.



Figure 7 Cut-off value comparison for absolute size criterion on input constraints on the DEC-LIN Track.

Figure 8 Cut-off value comparison for absolute size criterion on the DEC-LIN Track.

Figure 9 Cut-off value comparison for absolute size of maximum coefficient criterion on input constraints on the DEC-LIN Track. c = 1 is redundant since it would be equivalent to standard watched literal scheme.

Figure 10 Cut-off value comparison for absolute size of maximum coefficient criterion on the DEC-LIN Track. c = 1 is redundant since it would be equivalent to the standard watched literal scheme.

Figure 11 Cut-off value comparison for relative size criterion on input constraints on the DEC-LIN Track.

Figure 12 Cut-off value comparison for relative size criterion on the DEC-LIN Track.