



On Weakening Strategies for PB Solvers

Daniel Le Berre¹, Pierre Marquis^{1,2}, and Romain Wallon^{1(✉)}

¹ CRIL, Univ Artois & CNRS, Lens, France
{leberre,marquis,wallon}@cril.fr

² Institut Universitaire de France, Paris, France

Abstract. Current pseudo-Boolean solvers implement different variants of the cutting planes proof system to infer new constraints during conflict analysis. One of these variants is *generalized resolution*, which allows to infer strong constraints, but suffers from the growth of coefficients it generates while combining pseudo-Boolean constraints. Another variant consists in using *weakening* and *division*, which is more efficient in practice but may infer weaker constraints. In both cases, weakening is mandatory to derive conflicting constraints. However, its impact on the performance of pseudo-Boolean solvers has not been assessed so far. In this paper, new application strategies for this rule are studied, aiming to infer strong constraints with small coefficients. We implemented them in *Sat4j* and observed that each of them improves the runtime of the solver. While none of them performs better than the others on all benchmarks, applying weakening on the *conflict* side has surprising good performance, whereas applying *partial* weakening and division on both the conflict and the reason sides provides the best results overall.

Keywords: PB constraint · Constraint learning · Cutting planes

1 Introduction

The last decades have seen many improvements in SAT solving that are at the root of the success of modern SAT solvers [5, 13, 15]. Despite their practical efficiency on many real-world instances, these solvers suffer from the weakness of the resolution proof system they use in their conflict analyses. Specifically, when proving the unsatisfiability of an input formula requires an exponential number of resolution steps – as for pigeonhole-principle formulae [9] – a SAT solver cannot find a refutation proof efficiently. This motivated the development of pseudo-Boolean (PB) solvers [17], which take as input conjunctions of *PB constraints* (linear inequations over Boolean variables) and apply *cutting planes based inference* to derive inconsistency [8, 10, 16]. This inference system is stronger than the resolution proof system, as it *p-simulates* the latter: any resolution proof can be translated into a cutting planes proof of polynomial size [2]. Using such a proof system may, in theory, make solvers more efficient: for instance, a pigeonhole principle formula may be refuted with a linear number of cutting planes steps.

However, in practice, PB solvers fail to keep the promises of the theory. In particular, they only implement *subsets* of the cutting planes proof system, which degenerate to resolution when given a CNF formula as input: they do not exploit the full power of the cutting planes proof system [20]. One of these subsets is *generalized resolution* [10], which is implemented in many PB solvers [1, 4, 11, 19]. It consists in using the *cancellation* rule to combine constraints so as to resolve away literals during conflict analysis, as SAT solvers do with the resolution rule. Another approach has been introduced by *RoundingSat* [7], which relies on the *weakening* and *division* rules to infer constraints having smaller coefficients to be more efficient in practice. These proof systems are described in Sect. 2.

This paper follows the direction initiated by *RoundingSat* and investigates to what extent applying the *weakening* rule may have an impact on the performance of PB solvers. First, we show that applying a *partial* weakening instead of an aggressive weakening as proposed in [7] allows to infer stronger constraints while preserving the nice properties of *RoundingSat*. Second, we show that weakening operations can be extended to certain literals that are falsified by the current partial assignment to derive shorter constraints. Finally, we introduce a trade-off strategy, trying to get the best of both worlds. These new approaches are described in Sect. 3, and empirically evaluated in Sect. 4.

2 Pseudo-Boolean Solving

We consider a propositional setting defined on a finite set of classically interpreted propositional variables V . A *literal* l is a variable $v \in V$ or its negation \bar{v} . Boolean values are represented by the integers 1 (true) and 0 (false), so that $\bar{v} = 1 - v$. A *PB constraint* is an integral linear equation or inequation over Boolean variables. Such constraints are supposed, w.l.o.g., to be in the normalized form $\sum_{i=1}^n \alpha_i l_i \geq \delta$, where α_i (the *coefficients* or *weights*) and δ (the *degree*) are positive integers and l_i are literals. A *cardinality constraint* is a PB constraint with its weights equal to 1 and a *clause* is a cardinality constraint of degree 1.

Several approaches have been designed for solving PB problems. One of them consists in encoding the input into a CNF formula and let a SAT solver decide its satisfiability [6, 14, 18], while another one relies on lazily translating PB constraints into clauses during conflict analysis [21]. However, such solvers are based on the *resolution* proof system, which is somewhat *weak*: instances that are hard for resolution are out of reach of SAT solvers. In the following, we consider instead solvers based on the *cutting planes* proof system, the PB counterpart of the resolution proof system. Such solvers handle PB constraints natively, and are based on one of the two main subsets of cutting planes rules described below.

2.1 Generalized Resolution Based Solvers

Following the CDCL algorithm of SAT solvers, PB solvers based on *generalized resolution* [10] make *decisions* on variables, which force other literals to be satisfied. These *propagated* literals are detected using the *slack* of each constraint.

Definition 1 (slack). Given a partial assignment ρ , the slack of a constraint $\sum_{i=1}^n \alpha_i l_i \geq \delta$ is the value $\sum_{i=1, \rho(l_i) \neq 0}^n \alpha_i - \delta$.

Observation 1. Let s be the slack of the constraint $\sum_{i=1}^n \alpha_i l_i \geq \delta$ under some partial assignment. If $s < 0$, the constraint is currently falsified. Otherwise, the constraint requires all unassigned literals having a weight $\alpha > s$ to be satisfied.

Example 1. Let ρ be the partial assignment such that $\rho(a) = 1$, $\rho(c) = \rho(d) = \rho(e) = 0$ (all other variables are unassigned). Under ρ , the constraint $6\bar{b} + 6c + 4e + f + g + h \geq 7$ has slack 2. As \bar{b} is unassigned and has weight $6 > 2$, this literal is propagated (the constraint is the *reason* for \bar{b}). This propagation falsifies the constraint $5a + 4b + c + d \geq 6$, which now has slack -1 (this is a *conflict*).

When a conflict occurs, the solver analyzes this conflict to derive an *assertive* constraint, i.e., a constraint propagating some of its literals. To do so, it applies successively the *cancellation* rule between the conflict and the reason for the propagation of one of its literals (“LCM” denotes the least common multiple):

$$\frac{\alpha l + \sum_{i=1}^n \alpha_i l_i \geq \delta \quad \beta \bar{l} + \sum_{i=1}^n \beta_i l_i \geq \delta' \quad \mu\alpha = \nu\beta = \text{LCM}(\alpha, \beta)}{\sum_{i=1}^n (\mu\alpha_i + \nu\beta_i) l_i \geq \mu\delta + \nu\delta' - \mu\alpha} \text{ (canc.)}$$

To make sure that an assertive constraint will be eventually derived, the constraint produced by this operation has to be conflictual, which is not guaranteed by the cancellation rule. To preserve the conflict, one can take advantage of the fact that the slack is *subadditive*: the slack of the constraint obtained by applying the cancellation between two constraints is at most equal to the sum of the slacks of these constraints. Whenever the sum of both slacks is not negative, the constraint may not be conflictual, and the *weakening* and *saturation* rules are applied until the slack of the reason becomes low enough to ensure the conflict to be preserved (only literals that are not falsified may be weakened away).

$$\frac{\alpha l + \sum_{i=1}^n \alpha_i l_i \geq \delta}{\sum_{i=1}^n \alpha_i l_i \geq \delta - \alpha} \text{ (weakening)} \qquad \frac{\sum_{i=1}^n \alpha_i l_i \geq \delta}{\sum_{i=1}^n \min(\delta, \alpha_i) l_i \geq \delta} \text{ (saturation)}$$

Example 2 (Example 1 cont'd). As $5a + 4b + c + d \geq 6$ is conflicting and \bar{b} was propagated by $6\bar{b} + 6c + 4e + f + g + h \geq 7$, the cancellation rule must be applied between these two constraints to eliminate b . To do so, the *conflict side* (i.e., the first constraint) has to be multiplied by 3 and the *reason side* (i.e., the second constraint) by 2, giving slack -3 and 4, respectively. As the sum of these values is equal to 1, the resulting constraint is not guaranteed to be conflictual. Thus, the reason is weakened on g and h and saturated to get $5\bar{b} + 5c + 4e + f \geq 5$, which has slack 1. To cancel b out, this constraint is multiplied by 4 and the conflict by 5, giving $25a + 25c + 16e + 5d + 4f \geq 30$, which has slack -1 .

This approach has several drawbacks. Observe in Example 2 the growth of the coefficients in just one derivation step. In practice, there are many such steps during conflict analysis, and the learned constraints will be reused later on, so that coefficients will continue to grow, requiring the use of arbitrary precision arithmetic. Moreover, after each weakening operation, the LCM of the coefficients must be recomputed to estimate the slack, and other literals to be weakened must be found. The cost of these operations motivated the development of alternative proof systems, such as those weakening the derived constraints to infer only cardinality constraints [1], or those based on the *division* rule.

2.2 Division Based Solvers

To limit the growth of the coefficients during conflict analysis, *RoundingSat* [7] introduced an aggressive use of the weakening and *division* rules.

$$\frac{\sum_{i=1}^n \alpha_i l_i \geq \delta \quad r > 0}{\sum_{i=1}^n \lceil \frac{\alpha_i}{r} \rceil l_i \geq \lceil \frac{\delta}{r} \rceil} \text{ (division)}$$

When a conflict occurs, both the conflict and the reason are weakened so as to remove all literals not falsified by the current assignment and having a coefficient not divisible by the weight of the literal used as pivot for the cancellation, before being divided by this weight. This ensures that the pivot has a weight equal to 1, which guarantees that the result of the cancellation will be conflictual [3].

Example 3 (Example 2 cont'd). The weakening operation is applied on both the conflict $5a + 4b + c + d \geq 6$ and the reason $6\bar{b} + 6c + 4e + f + g + h \geq 7$, yielding $4b + c + d \geq 1$ and $6\bar{b} + 6c + 4e \geq 4$, respectively. Both constraints are then divided by the coefficient of the pivot (4 and 6, respectively), giving $b + c + d \geq 1$ and $\bar{b} + c + e \geq 1$. Applying the cancellation rule on these two constraints gives $2c + d + e \geq 1$, which is equivalent to the clause $c + d + e \geq 1$.

The *RoundingSat* approach succeeds in keeping coefficients small, and its aggressive weakening allows to find the literals to remove efficiently. However, some constraints inferred by this solver may be weaker than those inferred with generalized resolution (compare the constraints derived in Examples 2 and 3).

3 Weakening Strategies

As explained before, the weakening rule is mandatory in PB solvers to maintain the inferred constraints conflictual. In the following, we introduce different strategies for applying this rule in PB solvers, designed towards finding a tradeoff between the strength of the inferred constraints and their size.

3.1 Weakening Ineffective Literals for Shorter Constraints

Within CDCL solvers, one captures the reason for a conflict being encountered. A conflict occurs when a variable is propagated to both 0 and 1. Intuitively, understanding why such a conflict occurred amounts to understanding why these values have been propagated. In the PB case, a constraint may be conflicting (resp. propagate literals) even if it contains literals that are unassigned or already satisfied (see Example 1). However, conflicts (resp. propagations) depend only on *falsified* literals (the slack of a constraint changes only when one of its literals is falsified). Literals that are not falsified are thus *ineffective*: they do not play a role in the conflict (resp. propagation), and may thus be weakened away. We can go even further: when most literals are falsified, weakening some of them may still preserve the conflict (resp. propagation), as shown in the following example.

Example 4. Let ρ be the partial assignment such that $\rho(a) = \rho(c) = \rho(f) = 0$ (all other variables are unassigned). Under ρ , $3\bar{a} + 3\bar{b} + c + d + e \geq 6$ has slack 2, so that literal \bar{b} is propagated to 1. This propagation still holds after weakening away \bar{a} , d and e , giving after saturation $\bar{b} + c \geq 1$. Similarly, consider the conflicting constraint $2a + b + c + f \geq 2$. After the propagation of \bar{b} , weakening the constraint on c and applying saturation produces $a + b + f \geq 1$, which is still conflicting. In both cases, the slack allows to detect whether a literal can be weakened.

Observe that the constraints obtained are shorter, but are always clauses. This guarantees that the resulting constraint will be conflictual, but, if this operation is performed on both sides, only clauses can be inferred, and the proof system boils down to resolution, as in *SATIRE* [21] or *Sat4j-Resolution* [11].

Example 5 (Example 4 cont'd). If a resolution step is performed between the weaker constraints $\bar{b} + c \geq 1$ and $a + b + f \geq 1$, the clause $a + c + f \geq 1$ is inferred. However, if only one side is weakened, for example the conflict side, the cancellation between $3\bar{a} + 3\bar{b} + c + d + e \geq 6$ and $a + b + f \geq 1$ produces the constraint $3f + c + d + e \geq 3$. Observe that, when the weakening operation is applied at the next step, the stronger clause $c + f \geq 1$ is inferred after saturation.

3.2 Partial Weakening for Stronger Constraints

To avoid the inference of constraints that are too weak to preserve the strength of the proof system, an interesting option is to use *partial weakening*. Indeed, the weakening rule, as described above, can be generalized as follows.

$$\frac{\alpha l + \sum_{i=1}^n \alpha_i l_i \geq \delta \quad \varepsilon \in \mathbb{N} \quad 0 < \varepsilon \leq \alpha}{(\alpha - \varepsilon)l + \sum_{i=1}^n \alpha_i l_i \geq \delta - \varepsilon} \text{ (partial weakening)}$$

This rule is rarely used in practice by PB solvers, and the weakening rule (i.e., the case when $\varepsilon = \alpha$) is often preferred. However, partial weakening gives more freedom when it comes to inferring new constraints, and allows to infer stronger constraints. We implemented a variant of *RoundingSat* which uses this

rule as follows. Before cancelling a literal out during conflict analysis, all literals that are not currently falsified and have a coefficient not divisible by the weight of the pivot are *partially weakened* (instead of simply *weakened*). This operation is applied so that the resulting coefficient becomes a multiple of the weight of the pivot. This approach preserves the nice properties of *RoundingSat* (see [7, Proposition 3.1 and Corollary 3.2]), and in particular the fact that the produced constraint will be conflictual (the coefficient of the pivot will be equal to 1). It also allows to infer stronger constraints, as illustrated by the following example.

Example 6. Let ρ be the partial assignment defined by $\rho(a) = 1$ and $\rho(b) = \rho(c) = \rho(d) = \rho(e) = 0$ (all other variables are unassigned). Consider the (conflicting) constraint $8a + 7b + 7c + 2d + 2e + f \geq 11$, where b is the literal to be cancelled out. The above rule yields $7a + 7b + 7c + 2d + 2e \geq 9$ which, divided by 7, gives $a + b + c + d + e \geq 2$. This constraint is stronger than the clause $b + c + d + e \geq 1$ inferred by *RoundingSat*, which weakens away the literal a .

This variant has several advantages. First, its cost is comparable to that of *RoundingSat*: checking whether a coefficient is divisible by the weight of the pivot is computed with the remainder of the division of the former by the latter, which is the amount by which the literal must be partially weakened. Second, the constraints it infers may be stronger than that of *RoundingSat*. Yet, this strategy does not reduce the size of the constraints as much as the weakening of ineffective literals. To get the best of both worlds, we introduce tradeoff strategies.

3.3 Towards a Tradeoff

The previous sections showed that the weakening operation may help finding short explanations for conflicts, but may also infer weaker constraints. Several observations may guide us towards tradeoff applications of the weakening rule.

First, the key property motivating *RoundingSat* to round the coefficient of the pivot to 1 does not require it to be equal to 1 on *both* sides of the cancellation: actually, having a coefficient equal to 1 on only *one* side is enough to guarantee the resulting constraint to be conflicting [3]. Weakening only the reason or the conflict is thus enough to preserve this property, while maintaining coefficients low enough, as only one side of the cancellation may need to be multiplied.

Second, one may apply the weakening rule in a different manner to keep coefficients small so as to speed up arithmetic operations. A possible approach is the following, that we call *Multiply and Weaken*. Let r be the coefficient of the pivot used in the cancellation appearing in the reason and c that in the conflict. Find two values μ and ν such that $(\nu - 1) \cdot r < \mu \cdot c \leq \nu \cdot r$ (which can be done using Euclidean division). Then, multiply the reason by ν , and apply weakening operations on this constraint so as to reduce the coefficient of the pivot to $\mu \cdot c$. Note that, in order to preserve the propagation, this coefficient cannot be weakened directly. Instead, ineffective literals (as described above) are successively weakened away so that the saturation rule produces the expected reduction on the coefficient. Since this operation does not guarantee to preserve the conflict,

an additional weakening operation has to be performed, as for generalized resolution. Note that this approach may also derive clauses, even though this is not always the case, as shown in the following example.

Example 7. Let ρ be the partial assignment such that $\rho(a) = \rho(d) = 0$ and $\rho(e) = 1$ (all other variables are unassigned). Under ρ the constraint $5a + 5b + 3c + 2d + e \geq 6$ propagates b . The constraint $3\bar{b} + 2a + 2d + \bar{e} \geq 5$ becomes thus falsified. Instead of using the LCM of 3 and 5 (i.e., 15), the reason of b is weakened on e and partially on c to get, after saturation, $3a + 3b + 2d + c \geq 3$. The cancellation produces then $5a + 4d + c + \bar{e} \geq 5$.

4 Experimental Results

This section presents an empirical evaluation of the various strategies introduced in this paper. To make sure that their comparison only takes care of the underlying proof systems, and not of implementation details, we integrated all of them in *Sat4j* [11] (including *RoundingSat* proof system). The source code is available on *Sat4j* repository¹.

All experiments presented in this section have been run on a cluster equipped with quadcore bi-processors Intel XEON E5-5637 v4 (3.5 GHz) and 128 GB of memory. The time limit was set to 1200 seconds and the memory limit to 32 GB. The whole set of decision benchmarks containing “small” integers used in the PB competitions since the first edition [12] was considered as input.

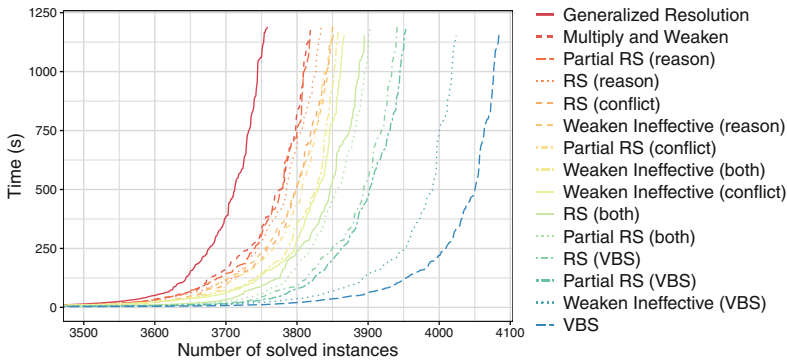


Fig. 1. Cactus plot of the different strategies implemented in *Sat4j*. For more readability, the first 3,500 easy instances are cut out.

As shown by Fig. 1, strategies applying heavily the weakening rule perform better than generalized resolution. Yet, among these strategies, none of them is strictly better than the others. In particular, the Virtual Best Solver (VBS),

¹ <https://gitlab.ow2.org/sat4j/sat4j/tree/weakening-investigations>.

obtained by choosing the best solver for each of the instances, performs clearly better than any individual strategy. Each of these individual strategies does not have an important contribution to the VBS, since the *both*, *conflict* and *reason* variants are very similar. However, if we consider the VBS of the different “main” strategies, and in particular that of *RS*, *Partial RS* and *Weaken Ineffective* (also represented on the cactus plot), their contributions become clearer: *Generalized Resolution* contributes to 6 instances, *RS* to 13 instances, *Partial RS* to 16 instances, and *Weaken Ineffective* to 83 instances. Even though *Multiply and Weaken* does not solve instances that are not solved by any other solver, it solves 13 instances more than 1 second faster than any other approach (5 among them are faster solved by more than 1 min). This suggests that choosing the *right* weakening strategy plays a key role in the performance of the solver.

The strategies showing the best and most robust performance are those applying the *RoundingSat* (*RS*) approach on both sides of the cancellation rule, as they widely take advantage of the inference power of the division rule. However, applying *partial* weakening instead of weakening gives better results, thanks to the stronger constraints it infers. In particular, *RS (both)* solves 3895 instances and *Partial RS (both)* solves 3903 instances (with 3845 common instances). The performance of *Partial RS (both)* is evidenced on the *tsp* family, especially on satisfiable instances: it solves 22 more such instances than *RS (both)*, i.e., 35 instances. For unsatisfiable instances, no common instances are solved: *Partial RS (both)* solves 7 instances, while *RS (both)* solves 5 distinct instances. In both cases, *Partial RS (both)* performs much more assignments per second than *RS (both)*, allowing it to go further in the search space within the time limit.

Surprisingly, another strategy exhibiting good performance consists in weakening ineffective literals on the *conflict* side at each cancellation (it contributes to 18 instances in the VBS). Similarly, *RoundingSat* strategies perform better when applied on the conflict side rather than the reason side. Since the early development of cutting planes based solvers, weakening has only been applied on the *reason* side (except for *RoundingSat* [7], which applies it on *both* sides). Our experiments show that it may be preferable to apply it only on the conflict side: literals introduced there when cancelling may still be weakened away later on.

The gain we observe between the different strategies has several plausible explanations. First, the solver does not explore the same search space from one strategy to another, and it may thus learn completely different constraints. In particular, they may be stronger or weaker, and this impacts the size of the proof built by the solver. Second, these constraints may be based on distinct literals, which may have side effects on the VSIDS heuristic [15]: different literals will be “bumped” during conflict analysis. Such side effects are hard to assess, due to the tight link between the heuristic and the other components of the solver.

5 Conclusion

In this paper, we introduced various strategies for applying the weakening rule in PB solvers. We showed that each of them may improve the runtime of the solver,

but not on all benchmarks. Contrary to the approaches implemented in most PB solvers, the strategies consisting in applying an aggressive weakening only on the conflict side provide surprisingly good results. However, approaches based on *RoundingSat* perform better, but our experiments showed that partial weakening is preferable in this context. This suggests that weakening operations should be guided by the strength of the constraints to infer. To do so, a perspective for further research consists in searching for better tradeoffs in this direction.

Acknowledgements. The authors are grateful to the anonymous reviewers for their numerous comments, that greatly helped to improve the paper. Part of this work was supported by the French Ministry for Higher Education and Research and the Hauts-de-France Regional Council through the “Contrat de Plan État Région (CPER) DATA”.

References

1. Chai, D., Kuehlmann, A.: A fast pseudo-boolean constraint solver. *IEEE Trans. CAD Integr. Circuits Syst.* **24**, 305–317 (2005)
2. Cook, W., Coullard, C.R., Turán, G.: On the complexity of cutting-plane proofs. *Discrete Appl. Math.* **18**, 25–38 (1987)
3. Dixon, H.: Automating pseudo-boolean inference within a DPLL framework. Ph.D. thesis, University of Oregon (2004)
4. Dixon, H.E., Ginsberg, M.L.: Inference methods for a pseudo-boolean satisfiability solver. In: *Proceedings of AAAI 2002*, pp. 635–640 (2002)
5. Eén, N., Sörensson, N.: An extensible SAT-solver. In: Giunchiglia, E., Tacchella, A. (eds.) *SAT 2003*. LNCS, vol. 2919, pp. 502–518. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24605-3_37
6. Een, N., Sörensson, N.: Translating pseudo-boolean constraints into SAT. *JSAT* **2**, 1–26 (2006)
7. Elffers, J., Nordström, J.: Divide and Conquer: towards faster pseudo-boolean solving. In: *Proceedings of IJCAI 2018*, pp. 1291–1299 (2018)
8. Gomory, R.E.: Outline of an algorithm for integer solutions to linear programs. *Bull. Am. Math. Soc.* **64**, 275–278 (1958)
9. Haken, A.: The intractability of resolution. *Theor. Comput. Sci.* **39**, 297–308 (1985)
10. Hooker, J.N.: Generalized resolution and cutting planes. *Ann. Oper. Res.* **12**, 217–239 (1988)
11. Le Berre, D., Parrain, A.: The SAT4J library, release 2.2, system description. *JSAT* **7**, 59–64 (2010)
12. Manquinho, V., Roussel, O.: The first evaluation of pseudo-boolean solvers (PB’05). *JSAT* **2**, 103–143 (2006)
13. Marques-Silva, J., Sakallah, K.A.: GRASP: a search algorithm for propositional satisfiability. *IEEE Trans. Comput.* **48**, 220–227 (1999)
14. Martins, R., Manquinho, V., Lynce, I.: Open-WBO: a modular MaxSAT solver. In: Sinz, C., Egly, U. (eds.) *SAT 2014*. LNCS, vol. 8561, pp. 438–445. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-09284-3_33
15. Moskewicz, M.W., Madigan, C.F., Zhao, Y., Zhang, L., Malik, S.: Chaff: engineering an efficient SAT solver. In: *Proceedings of DAC 2001*, pp. 530–535 (2001)
16. Nordström, J.: On the interplay between proof complexity and SAT solving. *ACM SIGLOG News* **2**, 19–44 (2015)

17. Roussel, O., Manquinho, V.M.: Pseudo-boolean and cardinality constraints. In: Handbook of Satisfiability, chap. 22, pp. 695–733. IOS Press (2009)
18. Sakai, M., Nabeshima, H.: Construction of an ROBDD for a PB-constraint in band form and related techniques for PB-solvers. IEICE Trans. Inf. Syst. **98**, 1121–1127 (2015)
19. Sheini, H.M., Sakallah, K.A.: Pueblo: a hybrid pseudo-boolean SAT solver. JSAT **2**, 165–189 (2006)
20. Vinyals, M., Elffers, J., Giráldez-Cru, J., Gocht, S., Nordström, J.: In between resolution and cutting planes: a study of proof systems for pseudo-boolean SAT solving. In: Beyersdorff, O., Wintersteiger, C.M. (eds.) SAT 2018. LNCS, vol. 10929, pp. 292–310. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-94144-8_18
21. Whitemore, J., Sakallah, J.K.K.A.: SATIRE: a new incremental satisfiability engine. In: Proceedings of DAC 2001, pp. 542–545 (2001)