

Efficient and Reliable Hitting-Set Computations for the Implicit Hitting Set Approach

Hannes Ihalainen¹ Dieter Vandesande^{2,3}, André Schidler⁴,
Jeremias Berg¹, Bart Bogaerts^{3,2}, Matti Järvisalo¹

¹University of Helsinki, Dept. Computer Science, Helsinki, Finland

²Vrije Universiteit Brussel, Dept. Computer Science, Brussels, Belgium

³KU Leuven, Dept. Computer Science, Leuven, Belgium

⁴University of Freiburg, Freiburg, Germany

hannes.ihalainen@helsinki.fi, dieter.vandesande@vub.be, schidler@cs.uni-freiburg.de, jeremias.berg@helsinki.fi,

bart.bogaerts@kuleuven.be, matti.jarvisalo@helsinki.fi

Abstract

The implicit hitting set (IHS) approach offers a general framework for solving computationally hard combinatorial optimization problems declaratively. IHS iterates between a decision oracle used for extracting sources of inconsistency and an optimizer for computing so-called hitting sets (HSs) over the accumulated sources of inconsistency. While the decision oracle is language-specific, the optimizer is usually instantiated through integer programming. We explore alternative algorithmic techniques for hitting set optimization based on different ways of employing pseudo-Boolean (PB) reasoning as well as stochastic local search. We extensively evaluate the practical feasibility of the alternatives in particular in the context of pseudo-Boolean (0–1 IP) optimization as one of the most recent instantiations of IHS. Highlighting a trade-off between efficiency and reliability, while a commercial IP solver turns out to remain the most effective way to instantiate HS computations, it can cause correctness issues due to numerical instability; in fact, we show that exact HS computations instantiated via PB reasoning can be made competitive with a numerically exact IP solver. Furthermore, the use of PB reasoning as a basis for HS computations allows for obtaining certificates for the correctness of IHS computations, generally applicable to any IHS instantiation in which reasoning in the declarative language at hand can be captured in the PB-based proof format we employ.

1 Introduction

The implicit hitting set (IHS) approach (Liffiton and Sakallah 2008) offers a general framework for solving computationally hard combinatorial optimization problems declaratively. Successful practical instantiations of IHS include solvers for maximum satisfiability (Davies and Bacchus 2011, 2013a,b; Davies 2014; Saikko, Berg, and Järvisalo 2016), finite-domain constraint optimization (Delisle and Bacchus 2013), answer set programming (Saikko et al. 2018), and pseudo-Boolean optimization (Smirnov, Berg, and Järvisalo 2021, 2022), as well as other computationally hard problems such as computing unsatisfiable subsets of formulas (Ignatiev et al. 2015; Gamba,

Bogaerts, and Guns 2023) and quantified Boolean formulas (Niskanen et al. 2022), propositional abduction (Ignatiev, Morgado, and Marques-Silva 2016; Saikko, Wallner, and Järvisalo 2016), and computing optimal causal graphs (Hyttinen, Saikko, and Järvisalo 2017).

Constituting a hybrid approach, the IHS framework consists of two main components: one for extracting sources of inconsistency in the declarative specification at hand (referred to as *core extraction*, with cores being constraints ruling out sources of inconsistency), and one for computing so-called hitting sets over sources of inconsistency (the *hitting set optimizer*). IHS iterates between these two components, accumulating cores and computing hitting sets over the so far accumulated cores. The hitting set obtained from the optimizer is used to focus the core extraction step (which reasons only on the original declarative specification) on finding sources of inconsistency that eliminate the current best candidate solution. In this way, the instance solved by the hitting set optimizer keeps on growing as more and more cores are extracted. The core extractor is implemented using a decision oracle specific to the declarative language at hand. In contrast, the optimizer is quite standardly instantiated through integer programming. Indeed, despite the fact that there has been significant early work put into investigating effective ways of realizing the hitting set component in IHS especially in the context of MaxSAT (Davies and Bacchus 2011, 2013a), the use of IP solvers remains largely the main mechanism used in modern IHS implementations for efficiency and for computing optimal solutions. Further improvements in the hitting set component are generally applicable to IHS implementations.

In this work we explore alternative algorithmic techniques for hitting set computations within IHS based, with the key aim of realizing effective *and* trustworthy hitting set computations. In terms of algorithmic alternatives to using IP solvers, we explore different ways of employing pseudo-Boolean (PB) reasoning (Roussel and Manquinho 2009) as well as stochastic local search (Chu, Cai, and Luo 2023; Chu et al. 2023) and their combinations together with IP solving. In terms of trustworthiness, it should be noted that the hitting set optimizers is in charge of ensuring that the objective function at hand is considered exactly. While IP solvers re-

main today the de facto approach to hitting set optimization within IHS, they are known to sometimes suffer from numerical instability, resulting in erroneous solutions (Cook et al. 2013). While numerically exact implementations of IP solvers exist (e.g., Eifler and Gleixner 2021; Cook et al. 2013), their performance remains today far from the efficiency of IP solvers in their default settings without guarantees on exactness.

The alternative hitting set computation approaches are applicable in general in different instantiations of the IHS approach, and—being based on exact reasoning—offer correctness guarantees. We extensively evaluate the practical feasibility of the alternative algorithmic approaches to hitting set optimization we propose in the context of pseudo-Boolean (0–1 IP) optimization as one of the most recent instantiations of IHS. Highlighting a trade-off between efficiency and reliability, while a commercial IP solver turns out to remain the most effective way to instantiate hitting computations in this setting, due to potential numerical instability using an off-the-shelf IP solver as-is is not generally exact; in fact, we show that exact HS computations instantiated via PB reasoning can be made competitive with a numerically-exact IP solver. Furthermore, the use of PB reasoning as a basis for hitting set computations allows for obtaining *certificates of correctness* of IHS computations. Indeed, while not yet generally adopted in IP solvers, the idea that solvers should be *certifying* (as in not only produce an answer but also a certificate of correctness of this answer), has slowly found its way from SAT solving to richer paradigms such as SMT solving (Barbosa et al. 2022), constraint programming (Gocht, McCreesh, and Nordström 2022; Flippo et al. 2024) and pseudo-Boolean optimization (see, e.g., (Roussel 2024)). Building on this, one of the key contributions of this paper is the development of the first *certifying instantiation* of a state-of-the-art IHS approach. To achieve this, we build on the VeriPB proof format (Bogaerts et al. 2023), which has recently been shown to be applicable to a rich variety of MaxSAT algorithms (Vandesande, De Wulf, and Bogaerts 2022; Berg et al. 2023, 2024; Jabs et al. 2025; Vandesande, Coll, and Bogaerts 2026), with IHS being the most important missing paradigm. Our results on certifying IHS are applicable to any IHS instantiation where the reason of the decision oracle can be captured with PB-based reasoning, and as such includes e.g. MaxSAT as well.

2 Preliminaries

2.1 Pseudo-Boolean Optimization

A *literal* ℓ is a Boolean variable x or its negation $\bar{x} = 1 - x$, where variables take values 0 (false) or 1 (true). A *pseudo-Boolean (PB) constraint* C is a 0–1 linear inequality $\sum_i a_i \ell_i \geq A$, where a_i and A are integers. Without loss of generality, we often assume PB constraints to be in *normal form*, meaning that the literals ℓ_i are over distinct variables, each *coefficient* a_i is *positive* and the *degree (of falsity)* A is non-negative. A *pseudo-Boolean formula* F is a conjunction $\bigwedge_j C_j$ of PB constraints. When convenient, we view F as a *set* of constraints. An *objective* O is an expression $\sum_i w_i \ell_i + lb$, where the coefficients w_i and the constant lb

are integers (lb stands for “lower bound”: when the w_i are all possible, this constant term is indeed a lower bound on the cost).

An assignment α maps variables to either 0 or 1. The assignment α satisfies a (normalized) PB constraint $C = \sum_i a_i \ell_i \geq A$ if $\sum_i a_i \alpha(\ell_i) \geq A$ and is a solution to a formula F if it satisfies all of its constraints. The cost $\alpha(O)$ of an assignment α under an objective $O = \sum_i w_i \ell_i + lb$ is $\sum_i w_i \alpha(\ell_i) + lb$ (we assume without loss of generality that all the w_i are positive). An instance of the *pseudo-Boolean optimization problem* is a tuple (F, O) , where F is a PB formula and O an objective under minimization. The solutions of (F, O) are the solutions of F . A solution α is optimal if $\alpha(O) \leq \beta(O)$ for each solution β of (F, O) . The optimal cost of (F, O) is $\alpha(O)$ for an optimal solution of (F, O) .

Following (Smirnov, Berg, and Järvisalo 2021), a constraint C is a *core* of (F, O) if the following conditions hold: (i) the literals in C are objective literals or their negations and (ii) all assignments α that satisfy F also satisfy C (denoted by $F \models C$). In other words, a core is an implied constraint that only mentions variables occurring in the objective.¹

The following proposition forms the basis for how the IHS approach to PBO employs cores to compute optimal solutions.

Proposition 1 (e.g., Smirnov, Berg, and Järvisalo 2021). *Let (F, O) be a PBO instance, α_{best} an optimal solution to (F, O) , \mathcal{K} a set of cores of (F, O) , and γ an optimal solution to (\mathcal{K}, O) . Then $\gamma(O) \leq \alpha_{best}(O)$.*

In words, Proposition 1 states that the cost of the optimal (minimum-cost) solutions to any set of cores is a lower bound on the optimal cost of the instance.

2.2 The Implicit Hitting Set Approach

Algorithm 1 details PBO-IHS, the implicit hitting set algorithm for pseudo-Boolean optimization (Smirnov, Berg, and Järvisalo 2021, 2022). The high-level overview of this algorithm (when ran on an instance (F, O)) is as follows. During the execution, collect a set \mathcal{K} of cores and repeatedly solve the optimization problem (\mathcal{K}, O) . When solving this problem to optimality, Proposition 1 guarantees that we obtain a lower bound on the optimal cost of (F, O) . Based on the outcome of such an optimization call, a decision oracle is then called under a set of assumptions (literals that are given a fixed value before search starts). As a result either a feasible solution (giving us an upper bound on the optimal cost) or a new core (that eliminates the previously found solution) is obtained. The iterations between the optimizer and the decision oracle are continued until the upper bound equals the lower bound (which is guaranteed to happen at the latest when all cores have been found).

In more detail, given a PBO instance (F, O) , PBO-IHS begins by checking the existence of solutions of (F, O) by invoking PB decision procedure PB-Solve on F (Line 2). The call returns an indicator *sat?* for satisfiability and a solution α_{best} in the positive case. Given α_{best} , an upper bound

¹See Section 2.3 for a historic explanation of this terminology.

```

1 PBO-IHS( $F, O$ )
   Input: A PBO instance  $(F, O)$ 
   Output: An optimal solution  $\alpha_{best}$ 
2  $(\alpha_{best}, sat?) \leftarrow \text{PB-Solve}(F)$ ;
3 if not sat? then
4   return “no feasible solutions”;
5  $ub \leftarrow \alpha_{best}(O)$ ;  $lb \leftarrow -\infty$ ;  $\mathcal{K} \leftarrow \text{seed}(F)$ ;
6 while TRUE do
7    $(\gamma, lb) \leftarrow \text{Solve-HS}(\mathcal{K}, O, lb, ub)$ ;
8   if  $ub = lb$  then break;
9    $(\mathcal{N}, \alpha) \leftarrow \text{Extract-Cores}(\gamma, F, O)$ ;
10  if  $\alpha(O) < ub$  then  $ub \leftarrow \alpha(O)$ ;  $\alpha_{best} \leftarrow \alpha$ ;
11  if  $ub = lb$  then break;
12   $\mathcal{K} \leftarrow \mathcal{K} \cup \mathcal{N}$ ;
13 return  $\alpha_{best}$ ;

```

Algorithm 1: IHS in the context of PBO.

ub and a lower bound lb on the optimal cost of the instance are initialized to $\alpha_{best}(O)$ and $-\infty$, respectively, on Line 5. During search, α_{best} will always contain the currently best-known solution which has the property that $\alpha_{best}(O) = ub$. A set \mathcal{K} of cores is initialized by *seeding* to be the set of all constraints in F that only mention objective literals. In other words, all constraints in F that are cores are immediately added to \mathcal{K} .

The main search loop (Lines 6-12) iterates until $lb = ub$, at which point α_{best} is known to be optimal. Each iteration begins with the hitting set¹ component of IHS, i.e., the procedure `Solve-HS` computing a solution γ of the instance (\mathcal{K}, O) consisting of the cores found so far and the objective (Line 7). The main focus in this work is on the `Solve-HS` procedure: we will detail various instantiations of `Solve-HS` in Section 3.

In addition to γ , each call to `Solve-HS` makes use of Proposition 1 to return a lower bound lb on the optimal cost of (F, O) . As detailed in Section 3, the instantiations we consider will return either the input parameter lb unchanged, or the cost $\gamma(O)$ when γ is known to be an optimal solution of (\mathcal{K}, O) . In a textbook implementation of plain IHS, `Solve-HS` would always search for an optimal solution of \mathcal{K} with respect to O . However, to make IHS perform well in practice, it is important to heuristically allow it to return suboptimal solutions as well. To guarantee termination, the heuristics must ensure a form of fairness: at any iteration of search, an optimal solution must always be returned at some later iteration. For more details and a formal proof of correctness of Algorithm 1, we refer the reader to (Bacchus et al. 2017).

Next, the procedure `Extract-Cores` is used to obtain new cores. This is done by repeated invocations of a decision procedure for PB constraints of F under a set of assumptions \mathcal{A} . This set of assumptions is initialized to the solution γ , meaning that first a solution of F that extends γ is sought for. If F is unsatisfiable under \mathcal{A} the decision oracle uses standard conflict analysis methods to extract a new core C of F that contradicts \mathcal{A} (a core that is false whenever all literals in \mathcal{A} are true) and adds it to the set \mathcal{N} of new

cores to be returned by `Extract-Cores` and added to \mathcal{K} in Line 12 of Algorithm 1. Before the next call to the decision procedure, the set \mathcal{A} is relaxed by the so-called *weight-aware core extraction* (Berg and Jarvisalo 2017) heuristic that—informally speaking—removes at least one of the literals in C from the assumptions, thus preventing C from being rediscovered in later iterations. The `Extract-Cores` procedure terminates when the decision procedure provides a solution α of $F \wedge \bigwedge_{\ell \in \mathcal{A}} \ell$. In case this solution has a better objective value than the currently stored best solution, the best solution is updated and termination is checked on Lines 10 and 11 of Algorithm 1.

2.3 A Remark on Terminology

Much of the terminology we use has a historic background, dating back to IHS approaches to MaxSAT (Davies and Bacchus 2011).

Firstly, our definition of “unsatisfiable core” differs from how the term is traditionally defined. Originally an “unsatisfiable core” was defined as a subset of the original formula that is unsatisfiable. In the context of assumption-based SAT solving (also employed for SAT-based optimization) fresh variables are used to represent constraints in the input formula and such a core became “a subset of the assumptions that cannot jointly be set to a specific value” (or alternatively, a clause over assumption variables learned by the solver). In the PBO setting, this was then further generalized to be any PB constraint over assumptions (objective variables) learned by the PB solver (Smirnov, Berg, and Jarvisalo 2021).

Secondly, this paper is about the so-called implicit hitting set approach. The notion of “hitting set” has as similar origin: when all cores are clauses over objective literals, the problem solved by `Solve-HS` is actually a minimum-cost hitting-set problem: in this case we search for a smallest-cost set of objective literals that hits each clause. In our more general PBO setting, however, `Solve-HS` itself needs to solve a PBO instance, but with the special property that the constraints only mention variables occurring in the objective.

In this paper, we opted to keep the historic terminology, even if it has become inaccurate due to the generality of the problem at hand. We do so to make the relationship with previous work on IHS for PBO (Smirnov, Berg, and Jarvisalo 2021) as well as for other domains more clear.

3 Alternatives Approaches to Hitting Set Computations

We now turn to investigating various alternative instantiations of the hitting set component (i.e., the `Solve-HS` procedure) for computing optimal solutions to (\mathcal{K}, O) . Currently efficient instantiations of `Solve-HS` relatively standardly make use of a commercial IP solver and floating-point computations. This holds true for the current state-of-the-art PBO-IHS solver, implementing IHS for PBO, as well as state-of-the-art IHS implementations for other declarative paradigms (Smirnov, Berg, and Jarvisalo 2021, 2022; Davies and Bacchus 2013a; Delisle and Bacchus 2013; Saikko et al. 2018; Gamba, Bogaerts, and Guns 2023; Saikko, Wallner, and Jarvisalo 2016; Hyttinen, Saikko, and Jarvisalo 2017).

```

1 Solve-HS( $\mathcal{K}, O, lb, ub$ )
2   if use-sls () then
3      $\gamma \leftarrow \text{Core-SLS}(\mathcal{K}, O)$ 
4     if  $\gamma(O) < ub$  then
5       return ( $\gamma, lb$ )
6   opt?  $\leftarrow \text{optimal-sol}()$ 
7    $\gamma \leftarrow \text{Optimize}(\mathcal{K}, O, ub, \text{opt?})$ ;
8   if  $\gamma(O) = ub$  or opt? then return ( $\gamma, \gamma(O)$ );
9   else return ( $\gamma, lb$ );

```

Algorithm 2: General view on the hitting set component, i.e., computing solutions to the cores extracted so far.

While numerically-stable IP solvers are available today, most IP solvers in their default settings employ floating-point computations for practical efficiency, essentially giving up guarantees on finding optimal solutions. This can and has caused intrinsic problems in terms of correctness of IHS. Towards even higher levels of correctness guarantees, verifiable proofs of optimality would be ideal. With these motivations, we explore alternative instantiations of Solve-HS based on recent developments in conflict-driven pseudo-Boolean solvers (Devriendt et al. 2021; Devriendt, Gleixner, and Nordström 2021) that allow producing certificates of correctness of the computations of Algorithm 1, as well as various ways of combining them within the hitting set component with IP solving and stochastic local search (SLS).

3.1 A General View on Hitting Set Solving

Algorithm 2 details a generalized abstraction of the Solve-HS procedure (the hitting set component), allowing for different ways of combining IP solving, conflict-driven pseudo-Boolean solving, and stochastic local search. PBO-IHS (Algorithm 1) is tasked with computing an optimal solution to an instance (F, O) . It will invoke Solve-HS with the set \mathcal{K} of cores extracted so far, the objective O , and the current upper and lower bound (ub, lb) on the optimal cost of (F, O) .

First Solve-HS invokes use-sls, a user-specified heuristic for determining whether SLS should be invoked on Line 2. If the solution γ that SLS returns has a cost lower than the current upper bound ub , Solve-HS terminates on Line 5. Otherwise, the subroutine Optimize is invoked on Line 7 and tasked to either compute a solution γ of (\mathcal{K}, O) that has $\gamma(O) < ub$, or (if no such solutions exist) an optimal solution γ of (\mathcal{K}, O) . Since the current best solution α_{best} of (F, O) maintained by PBO-IHS has $\alpha_{best}(O) = ub$ and is also a solution of (\mathcal{K}, O) , an alternative view on Optimize is that of either finding a solution of (\mathcal{K}, O) of cost lower than ub or determining that such a solution does not exist. As an additional heuristic, we allow forcing Optimize to compute an optimal solution to (\mathcal{K}, O) , essentially forcing an improvement to the lower bound. More specifically, the user-specified heuristic optimal-sol invoked on Line 6 returns true if Optimize is only allowed to terminate with an optimal solution to (\mathcal{K}, O) .

Before terminating, Solve-HS checks whether the lower

bound can be improved to $\gamma(O)$ on Line 8. This can happen if either Optimize was forced to compute an optimal solution or if it was unable to find a solution that has a cost lower than ub . Otherwise, the procedure terminates and returns γ and the input bound lb on Line 9.

Next, we detail different instantiations of Optimize that we consider as well as integration of local search for potential practical runtime improvements. Our general view on computing hitting sets allows for combining these instantiations in various different ways.

Optimize using an IP Solver. The arguably most commonly used instantiation of Optimize invokes a commercial IP solver on (\mathcal{K}, O) . We consider two variants of IP solving for instantiating Optimize. The first variant employs fast floating-point arithmetic in its computation. These are arguably among the most efficient methods of computing solutions (\mathcal{K}, O) , and most state-of-the-art IHS approaches instantiate Solve-HS with a floating-point IP solver. One disadvantage of such solvers, however, is that there have been repeated reports of such solvers outputting faulty solutions (Cook et al. 2013), preventing their use in applications that require trust in the results. Thus, we also consider exact IP solving, i.e., IP solving methods that employ exact arithmetic to circumvent issues related to floating point precision and rounding errors (Cook et al. 2013; Eifler and Gleixner 2023). This increased reliability of the results does, however, come at a cost in terms of computational efficiency. Recent computational studies (Eifler and Gleixner 2023) report slowdowns of a factor 8.1 when comparing a state-of-the-art exact IP solver to its floating point counterpart.

Optimize using Conflict-Driven PB Optimization. Complementing the branch & cut search commonly employed in IP solvers, we also consider harnessing recent developments in native pseudo-Boolean optimization procedures that reduce the problem of computing an optimal solution of (\mathcal{K}, O) into a sequence of decision problems. The decision problems are tackled with a decision procedure similar to the one employed within Algorithm 1 to extract cores (Devriendt et al. 2021; Devriendt, Gleixner, and Nordström 2021). More specifically, we consider three different variants of pseudo-Boolean optimization algorithms for computing an optimal solution to (\mathcal{K}, O) : solution-improving search (SIS), core-guided search (CG), and core-boosted search (CB).

A SIS-based approach maintains an upper bound ub_c on the optimal cost of (\mathcal{K}, O) . In the context of Solve-HS, ub_c is instantiated to the input upper bound ub . In each iteration, the decision procedure is invoked on $\mathcal{K} \wedge (O < ub_c)$, i.e., the cores in \mathcal{K} and a *solution-improving constraint* $(O < ub_c)$. If a solution exists, the returned solution γ will have $\gamma(O) < ub_c$, so the upper bound is updated, and the process repeats. Otherwise (if the decision procedure reports unsatisfiable), the latest solution found is determined to be optimal for (\mathcal{K}, O) .

One potential issue with this SIS-based approach is that the constraints $O < ub_c$ that are added to the solver are not sound for future calls to this oracle. Indeed, when new cores are learned, ub_c might no longer be an upper bound on the

best objective value. As a consequence, a fresh instantiation of the optimization solver needs to be created each time it is called. To deal with this, we also consider a variant that we call reified SIS. An incremental SIS approach adds the reified solution-improving constraint

$$r_{ub_c} \Rightarrow (\gamma(O) < ub_c),$$

i.e., the linear inequality

$$-Mr_{ub_c} + \gamma(O) < ub_c$$

where r_{ub_c} is a fresh variable and M is a large enough constant, and invokes the decision procedure while assuming $r = 1$. The benefit of relaxing the solution-improving constraint is that the decision procedure instantiation can be incrementally maintained between iterations of `Solve-HS`. Not just the constraint itself, but also all constraints derived from it will be reactivated in all future solve calls where a bound of ub_c or lower is enforced.

A core-guided approach (Devriendt et al. 2021) maintains a reformulated objective O^R , initialized to O , and a set \mathcal{C} of additional constraints, initialized to \emptyset . In each iteration, the decision procedure is queried for a solution of $\mathcal{K} \wedge \mathcal{C}$ while assuming all literals in O^R to 0. More informally, the query asks for a solution γ of $\mathcal{K} \wedge \mathcal{C}$ that has $\gamma(O^R) = 0$. If a solution exists, the obtained solution is an optimal solution of (\mathcal{K}, O) and the procedure terminates. Otherwise, the call returns a new core falsified by the assumptions. Next, the core-guided approach introduces new constraints to \mathcal{C} and reformulates O^R in a manner that resolves the inconsistency represented by the core in a cost-minimal way. Informally speaking, the invariant maintained by a core-guided search is that optimal solutions of $(\mathcal{K} \wedge \mathcal{C}, O^R)$ are optimal solutions of (\mathcal{K}, O) . Since the set \mathcal{K} of cores that we invoke core-guided search on monotonically increases during search, the state of the core-guided solver is maintained between invocations. While such *incremental invocations* of core-guided search prevent fixing any variables based on the bounds on the optimal cost of (\mathcal{K}, O) , variables are still fixed based on the bounds of the full instance (F, O) . For a concrete example, we fix the value of an objective literal ℓ that has a constant larger than ub to 0 since any solution that assigns $\ell = 1$ would incur a cost larger than ub .

Finally, core-boosted search (Berg, Demirovic, and Stuckey 2019) is a hybrid of core-guided and solution-improving search that starts with core-guided search for a user-specified number of iterations. After this, solution-improving search is invoked on the final reformulated instance $(\mathcal{K} \wedge \mathcal{C}, O^R)$ obtained from the core-guided search phase.

It is important to note that SIS, CG and CB can each provide intermediate solutions that allow terminating search as soon as a solution of cost lower than ub is found. Specifically, the instantiation of core-guided search that we use makes use of the so-called *stratification* heuristic (Ansótegui et al. 2013) and *weight-aware core extraction* (Berg and Jarvisalo 2017) that provide intermediate solutions, whereas plain core-guided search does not find any intermediate solutions before finding the first solution which is guaranteed to be optimal.

3.2 Integrating Stochastic Local Search

Stochastic local search (SLS) is a lightweight heuristic approach to pseudo-Boolean optimization that quickly finds low-cost solutions. In contrast to the approaches described in Section 3.1, SLS alone can not give any guarantee on optimality of a solution and cannot therefore establish a lower bound on its own. Hence an IHS approach with SLS integration in the hitting set component cannot terminate without an instantiation of `Optimize` that can check whether the solution found by SLS is optimal.

Our SLS component of hitting set computations (the `Core-SLS` procedure in Algorithm 2) is designed to achieve two desirable properties of SLS search in the context of `Solve-HS`. First, the search should incur as little overhead as possible to justify replacing calls to `Optimize` that can result in improved bounds on the optimal cost. Second, the solutions found by SLS should, intuitively, be as helpful as those found by `Optimize` in guiding IHS search toward an optimal solution of the full instance.

Our instantiation of `Core-SLS` is built on recent advances in SLS for pseudo-Boolean optimization (Chu et al. 2023) and adapts the NuPBO approach to the hitting set context. NuPBO starts from a complete assignment that is not necessarily a solution. In each iteration, the search flips the value of one variable in the hopes of eventually finding a solution. The choice of variable is based on a score that considers three aspects: (i) the contribution of the current variable value to satisfying the cores, (ii) how much flipping the variable value would contribute towards satisfying the cores, and (iii) how flipping the variable value would impact the cost of the assignment.

In more detail, we design `Core-SLS` for incremental calls, to run in conjunction with another instantiation of `Optimize`, and to produce good solutions for core extraction. Adapting the SLS approach to incremental calls is straightforward, as we can keep the state from the previous call and simply add the new constraints representing additional cores. Effectively, `Core-SLS` seeks to repair the previous solution to satisfy the new constraints. and (iii) how flipping the variable value would impact the cost of the assignment.

The solution returned by `Solve-HS` is used to extract the next set of cores. Thus it is important that the solutions returned by `Solve-HS` are not only of low cost, but also guide the IHS solver towards an optimal solution. We incorporate this idea in our SLS integration by ensuring that the solutions are *diversified*, i.e., that the subsequent solutions returned by `Solve-HS` have a large Hamming-distance, whenever possible. Hence, we split `Core-SLS` into two phases: (i) we start from the previous solution and (ii) we restart the search and start from the previous solution but randomly flip a large number of variable values. Whenever we find a solution in both phases, we prefer the one of lower cost. In case of a tie, we prefer the solution from the second phase to foster diversification. Further, `Core-SLS`'s purpose is avoiding costly calls to the optimizer. Hence we skip `Core-SLS` whenever we know from prior iterations that the optimizer calls do not take longer than `Core-SLS` calls.

4 Certified IHS Computations

We now turn to developing *proofs*, i.e., independently checkable certificates of correctness of results, for the IHS paradigm. We employ the VERIPB format as a general-purpose proof format for pseudo-Boolean optimization, combining implicational reasoning using cutting planes operations (such as adding up two PB constraints) (Cook, Coullard, and Turán 1987) with well-chosen strengthening rules that allow deriving non-implied constraints (Bogaerts et al. 2023). For representational succinctness of the current paper, details of the actual proof rules are not critical to spell out, but it is relevant that the strengthening rules allow us to derive fresh reification variables. In other words, if x is a variable that does not yet appear in the proof and C is a PB constraint, we can always derive (one or both) two new PB constraints $x \Rightarrow C$ and $x \Leftarrow C$ that express that x is true precisely when C is satisfied. If C is the normalized constraint $\sum_i a_i \ell_i \geq A$, then $x \Rightarrow C$ and $x \Leftarrow C$ are

$$A\bar{x} + \sum_i a_i \ell_i \geq A \text{ and} \\ (\sum_i a_i - A + 1)x + \sum_i a_i \bar{\ell}_i \geq \sum_i a_i - A + 1,$$

respectively.

4.1 VeriPB Proofs for PBO-IHS

Pseudo-Boolean solvers using solution-improving search and core guided search have had support for VERIPB proof logging for a while (see, e.g., the recent certifying track in the PBO Competitions (Roussel 2024)). As such, a natural idea toward obtaining uniform proofs for PBO-IHS is to instantiate SOLVE-HS with a PB-based approach and to combine the proofs produced by the SOLVE-HS subroutine with the proofs produced by the EXTRACT-CORES subroutine. While the idea is conceptually simple, there are a number of details that need to be considered. Due to space restrictions, we do not go into detail about how the proof production for the different components works. Instead, we discuss challenges in their integration and how they are addressed.

Constraint ID and Variable Management. With two different solvers co-producing parts of the proof, neither of the solvers are fully aware of which constraints are derived in the proofs and, in case they introduce new variables, which variables are unused so far. To facilitate this, we developed a general pseudo-Boolean proof logging API which both solvers use and which is responsible for maintaining relevant information, such as mappings from solver-specific variables to their name used in the proof. We expect this API to be useful more generally for future pseudo-Boolean proof logging approaches as well.

Non-Reified Solution-Improving Constraints. Another issue that needs to be addressed is that the optimizer uses (SOLVE-HS) is not optimizing the original problem, but a derived problem. When using a pseudo-Boolean optimizer with proof logging, it will log solutions it finds and might then use the associated solution-improving constraint expresses that from now on we are only interested in better solutions. Taking a holistic perspective, however, this solution-improving constraint is not sound: the solution SOLVE-HS

finds is not a solution of the whole problem. We already described in Section 3 two versions of integrating a SIS solver (incremental and non-incremental). In the incremental case the solver does not introduce a proper solution-improving constraint but a reified version thereof; this is easy to deal with from a proof logging perspective. In the non-incremental version, however, the solver is ran without assumptions and restarted upon every call. In this case the solver will be deriving those unsound solution-improving constraints. Our solution for dealing with this is to allow for a small discrepancy between the constraints the solver sees and the constraints derived in the proof, ensuring that the solution-improving constraint logged in the proof is indeed reified.

Example 2. Assume that the objective at hand is $x_1 + 2x_2 + 3x_3$ and during an execution of SOLVE-HS we find a solution with $x_1 = x_2 = 1$ and $x_3 = 0$. While this is a solution of \mathcal{K} , we do not know if this can be extended to a solution of F . Deriving the solution-improving constraint

$$x_1 + 2x_2 + 3x_3 < 3 \quad (1)$$

$$\text{(which normalizes to } \bar{x}_1 + 2\bar{x}_2 + 3\bar{x}_3 \geq 4) \quad (2)$$

in the proof would be unsound. In reified SIS this is not a problem since this constraint is not learned unconditionally. In regular SIS, however, when adding proof logging, we do not wish to change the behaviour of the solver. To this end, we add the above constraint to the solver but in the proof instead derive the reification constraint

$$4\bar{r}_3 + \bar{x}_1 + 2\bar{x}_2 + 3\bar{x}_3 \geq 4 \quad (3)$$

Later on, whenever the solver uses constraint (2) for deriving new constraints internally, the proof will instead reference constraint (3), resulting in all derivations of it to be conditional on r_3 . If it later turns out that there is solution of cost 3 for the original PBO instance, we can derive without loss of generality that r_3 is indeed true.

4.2 Proofs and IP Instantiations of SOLVE-HS

To harness both the effectiveness of SLS and floating-point IP solvers as well as proof-production capabilities of pseudo-Boolean instantiations of SOLVE-HS, we outline hybrid instantiations of OPTIMIZE that combine these strengths. The intuition here is that no proofs are needed for the calls to SOLVE-HS that do not result in optimal solutions of (\mathcal{K}, O) . The hybrid schemes we consider seek to minimize the calls to the (less efficient) proof-producing optimizers, while still enabling proofs for the overall IHS algorithm. More specifically, we consider the following hybridizations of non-proof-producing (i.e., SLS or IP) instantiations of OPTIMIZE, and the proof-producing ones.

Verify optimal solution (OptLB). The first approach invokes a proof-producing optimizer whenever the lower bound would be increased to equal the upper bound, i.e., when an optimal solution to the overall instance has been found. Intuitively, this represents the minimum amount of exact computation in SOLVE-HS needed to obtain a proof of correctness for the overall IHS algorithm. However, no

guarantees on the intermediate lower bounds are obtained. Additionally, re-invoking a proof-producing optimizer when the LB would be increased might require some amount of duplicate work in the optimizer recomputing a solution already found by the IP or SLS solver.

Verify all lower bounds (AII LB). The second approach uses a proof-producing optimizer whenever the lower bound would be refined, i.e., whenever the if statement on Line 8 of Algorithm 2 is true. Verifying all of the lower bounds requires more (potentially duplicated) exact computation. Nevertheless, this allows for generating proofs for more search heuristics used by instantiations of IHS that we are aware of.

Switch when forced lower bounds (ForceLB). The previous two hybridizations possibly perform duplicate work (asking a proof-producing optimizer to confirm all/some optimality claims of other optimizers), which might be undesirable. In this third hybridization, we make an a priori distinction between calls for which optimality is required and calls for which optimality is not required. Specifically, this version uses proof-producing optimizers to *only* instantiate the calls to `Optimize` in which the parameter `opt?` is true. The calls to `Optimize` that are not required to find an optimal solution can be instantiated with non-proof producing optimizers. This ensures that any calls to `SolveHS` that are required to find an optimal solution are indeed proof-producing. In this setting, the case $\gamma(O) = ub$ of the if statement on Line 8 that represents `Optimize` returning an optimal solution even when `opt?` is false, is not used since non-proof-producing optimizers may make such calls.

5 Experiments

In our experiments, we evaluate the performance of the different instantiations of `Optimize` as well as the overhead caused by proof logging. First, we compare how different solvers used as the single solver for `Optimize` perform as well as the hybrid configurations outlined in Section 4. Second, we evaluate the benefits of integrating SLS. Finally, we investigate the impact of proof logging on runtime performance. All implementation code, together with run scripts and raw data, can be found online (Ihalainen et al. 2025).

Setup The experiments reported on were run on compute nodes with two AMD EPYC 7513 CPUs, each with 64 cores running at 2.6 GHz. We reduced runtime variance by reserving whole CPUs for our experiments and run 16 instances in parallel on one CPU. Each run is limited to 14 GB memory and 1 h runtime using `runlim`², and each solver is executed once on each instance.

Solvers All our implementation builds on the PBHS solver.³ We use Gurobi 11.0.3 and SCIP 9.2.3⁴ (Bolusani et al. 2024) as IP solvers. The latter offers an exact mode (see

²<https://github.com/arminbiere/runlim>

³<https://bitbucket.org/coreo-group/pbhs>

⁴<https://github.com/scipop/scip> (commit 3045f20)

Solver	Single		Hybrid					
	#	m	AII LB		ForceLB		OptLB	
	#	m	#	m	#	m	#	m
Gurobi	951	0.66	-	-	-	-	-	-
SCIP	913	1.16	-	-	-	-	-	-
Exact SCIP	762	2.10	-	-	-	-	-	-
Core Guided	687	0.63	730	0.74	725	0.72	729	0.73
SIS Reified	651	1.52	654	1.41	648	1.59	704	1.22
SIS	672	0.79	684	0.86	697	0.88	696	0.80
SIS + CB	716	2.62	737	1.55	731	2.08	746	1.46

Table 1: Different instantiations of `Optimize`. *Single* shows instantiations using only a single solver and *Hybrid* shows instantiations that combine an IP solver (Gurobi) with a proof producing solver (Roundingsat). #: number of instances solved, m: average memory usage in GB.

Section 3.1). Further, we use a modified version of `Roundingsat`⁵ (Gocht and Nordström 2021; Bogaerts et al. 2023) as the proof producing PB solver and `VeriPB 2.2.2`⁶ to check these proofs. Proof logging is enabled only for experiments concerning proof logging.

Instances We evaluate the solvers on the same set of 1786 instances across tens of different benchmark domains as used in the the original papers presenting PBO-IHS (Smirnov, Berg, and Jarvisalo 2021, 2022); see (Smirnov 2021) for details.

5.1 Impact of Hitting Set Oracles

In the first experiment, we use a single solver for the hitting set computation in `Optimize`. Table 1 shows the results for different solvers, split into IP solvers (top) and designated PB solving methods. These results show that the IP solvers solve more instances. However, the difference between SCIP and Exact SCIP show that guaranteed optimality causes a significant performance decrease. Indeed, the number of solved instances decreases by more than 15% while the memory usage doubles.

As far as the dedicated PBO methods are concerned, first we notice that the variant of solution-improving search with reified solution-improving constraints appears to suffer from too much memory usage and is outperformed by pure SIS. This difference disappears when using the OptLB hybrid strategy and the number of calls to the PBO oracle is significantly decreased. Over the line, we observe that the combination of solution-improving search with core-boosting outperforms the SIS as well as the core-guided implementation.

The dedicated PBO methods solve fewer instances than the IP solvers, but they do come with certification, motivating the hybrid approaches. The more we delay the call to `Roundingsat`, the better the results: AII LB has the most overhead and performs overall poorest, while OptLB requires the fewest `Roundingsat` calls and performs best overall.

⁵<https://gitlab.com/MIAOresearch/software/roundingsat> (Commit: c548e109)

⁶<https://gitlab.com/MIAOresearch/software/VeriPB>

5.2 Impact of Integrating Local Search

Secondly, we integrated `Core-SLS` (cf. Section 3.2) into the single solver instantiations of `Optimize` and the best hybrid approach `OptLB`. By Table 2, integration of SLS improves the runtime performance of each of the instantiations. We use two SLS configurations, one for SLS combined with an IP solver and one for SLS combined with dedicated PBO approaches. In the hybrid configuration, we use the SLS configuration for IP. Integration of SLS consistently results in more instances solved, decrease memory usage and improve runtimes. This is most striking for `Exact SCIP`, for which SLS gives the biggest gain in the number of solved instances, also decreasing the memory consumption by almost 20%. Interestingly, often large numbers of instances are solved either with or without SLS but not both (Column d in Table 2). However, when SLS does not find a sufficiently low-cost solution it directly incurs a runtime overhead. Despite diversification, SLS can get stuck in local optima and cannot escape it without an optimizer. This suggests that improving SLS heuristics further for the IHS use case has potential for even further improvements.

In the hybrids, SLS can only replace IP solver calls due to inexactness. Here, SLS contributions significantly decrease, as the number of instances that SLS performs well on is balanced with instances it does poorly on. It is unlikely that instances the IP solver can only solve when SLS is used could be solved by the dedicated PBO approach.

5.3 Proof Logging Overhead

We evaluate proof logging using the most performant proof-producing variant, `OptLB` using `SIS + CB`. The proofs are checked using `VeriPB` with a timelimit of 10 h. By proof logging, we found several occasions where the IP solver wrongly claimed optimality; Gurobi on four instances a total of nine times, and four wrong results for `Exact SCIP`.

`SIS + CB` solved 746 instances without proof logging and 741 with proof logging. The median runtime overhead caused by proof logging is very low, 10.7%, while the overhead at the 90th percentile is 100.0% (see Figure 1). Although beyond the scope of this paper’s contributions, with

Solver	IHS + SLS			OptLB + SLS		
	#	m	d	#	m	d
Gurobi	+17	-0.02	31	-	-	-
SCIP	+4	-0.14	18	-	-	-
Exact SCIP	+44	-0.48	58	-	-	-
Core Guided	+41	+0.02	53	+11	-0.02	21
SIS Reified	+10	-0.06	52	+7	-0.04	31
SIS	+13	+0.02	39	+7	-0.02	25
SIS + CB	+19	-0.53	53	+3	-0.03	27

Table 2: Different instantiations of `Optimize` with and without SLS. #: number of solved instances (#); m: average memory usage in GB (positive numbers indicate an increase with SLS); d: number of instances solved with or without SLS but not by both configurations.

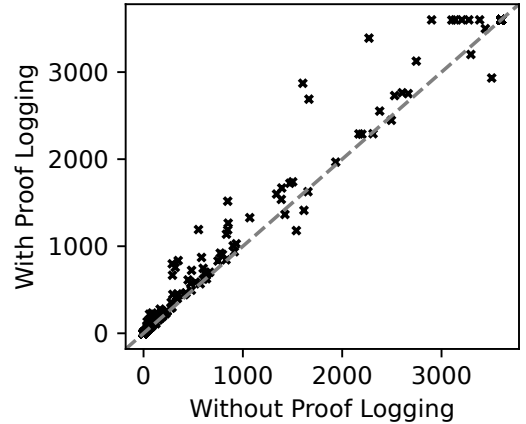


Figure 1: Comparison of runtime with and without proof logging. Times are in seconds.

the present version of the `VeriPB` proof checker, the median proof *checking* time is 588.0% of solving with proof logging; 712 proofs could be checked with `VeriPB` with a 10 h time limit. We expect these numbers to decrease significantly with on-going advances in `VeriPB` proof checking (Oertel 2025) once extended to optimization.

6 Conclusion

The hitting set component in IHS is standardly realized using (typically commercial) IP solvers using floating-point computations. However, this can lead to correctness issues in practice due to numerical issues and at least at present hinders gaining trustworthy IHS computations due to difficulties in uniformly certifying the computations of both the core extraction and the hitting set components. With these motivations, focusing on IHS for PBO as one of the most recent successful IHS instantiations, we explored different combinations of recent advances in exact PB reasoning, local search and IP solving for instantiating the hitting set component as potential means of gaining trustworthiness to IHS computations, addressing both issues with numerical instability and the challenge of realizing certified IHS. We extensively evaluate a wide range of different variants of such combinations, and showed that by clever combinations of IP solving, PB reasoning and local search, certified IHS can be realized and uniform proofs obtained for IHS computations while remaining reasonably performant, especially when compared to employing numerically-exact IP solving. The ideas presented are expected to be applicable to various other IHS instantiations, such as state-of-the-art IHS-based solvers for `MaxSAT`, which is a promising direction for further work.

Acknowledgments

This work is partially funded by the European Union (ERC, CertiFOX, 101122653). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

In addition, this work is partially funded by the Fonds Wetenschappelijk Onderzoek – Vlaanderen (G064925N and G070521N & fellowship 11A5J25N) and by the Research Council of Finland (grants 362987 and 356046), and by an Amazon Research Award (Fall/2023).

The authors acknowledge support by the state of Baden-Württemberg through bwHPC and the German Research Foundation (DFG) through grant INST 35/1597-1 FUGG.

References

- Ansótegui, C.; Bonet, M. L.; Gabàs, J.; and Levy, J. 2013. Improving WPM2 for (Weighted) Partial MaxSAT. In *CP*, volume 8124 of *Lecture Notes in Computer Science*, 117–132. Springer.
- Bacchus, F.; Hyttinen, A.; Jarvisalo, M.; and Saikko, P. 2017. Reduced Cost Fixing in MaxSAT. In *CP*, volume 10416 of *LNCS*, 641–651. Springer.
- Barbosa, H.; Reynolds, A.; Kremer, G.; Lachnitt, H.; Niemetz, A.; Nötzli, A.; Ozdemir, A.; Preiner, M.; Viswanathan, A.; Viteri, S.; Zohar, Y.; Tinelli, C.; and Barrett, C. W. 2022. Flexible Proof Production in an Industrial-Strength SMT Solver. In *IJCAR*, volume 13385 of *LNCS*, 15–35. Springer.
- Berg, J.; Bogaerts, B.; Nordström, J.; Oertel, A.; Paxian, T.; and Vandesande, D. 2024. Certifying Without Loss of Generality Reasoning in Solution-Improving Maximum Satisfiability. In *CP*, volume 307 of *LIPICs*, 4:1–4:28. Schloss Dagstuhl.
- Berg, J.; Bogaerts, B.; Nordström, J.; Oertel, A.; and Vandesande, D. 2023. Certified Core-Guided MaxSAT Solving. In *CADE*, volume 14132 of *LNCS*, 1–22. Springer.
- Berg, J.; Demirovic, E.; and Stuckey, P. J. 2019. Core-Boosted Linear Search for Incomplete MaxSAT. In *CPAIOR*, volume 11494 of *LNCS*, 39–56. Springer.
- Berg, J.; and Jarvisalo, M. 2017. Weight-Aware Core Extraction in SAT-Based MaxSAT Solving. In *CP*, volume 10416 of *Lecture Notes in Computer Science*, 652–670. Springer.
- Bogaerts, B.; Gocht, S.; McCreesh, C.; and Nordström, J. 2023. Certified Dominance and Symmetry Breaking for Combinatorial Optimisation. *J. Artif. Intell. Res.*, 77: 1539–1589.
- Bolusani, S.; Besançon, M.; Bestuzheva, K.; Chmiela, A.; Dionísio, J.; Donkiewicz, T.; van Doornmalen, J.; Eiffler, L.; Ghannam, M.; Gleixner, A.; Graczyk, C.; Halbig, K.; Hedtke, I.; Hoen, A.; Hojny, C.; van der Hulst, R.; Kamp, D.; Koch, T.; Kofler, K.; Lentz, J.; Manns, J.; Mexi, G.; Mühmer, E.; Pfetsch, M. E.; Schlösser, F.; Serrano, F.; Shinano, Y.; Turner, M.; Vigerske, S.; Weninger, D.; and Xu, L. 2024. The SCIP Optimization Suite 9.0. Technical report, Optimization Online.
- Chu, Y.; Cai, S.; and Luo, C. 2023. NuWLS: Improving Local Search for (Weighted) Partial MaxSAT by New Weighting Techniques. In *AAAI*, 3915–3923. AAAI Press.
- Chu, Y.; Cai, S.; Luo, C.; Lei, Z.; and Peng, C. 2023. Towards More Efficient Local Search for Pseudo-Boolean Optimization. In Yap, R. H. C., ed., *29th International Conference on Principles and Practice of Constraint Programming, CP 2023, August 27-31, 2023, Toronto, Canada*, volume 280 of *LIPICs*, 12:1–12:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Cook, W. J.; Coullard, C. R.; and Turán, G. 1987. On the complexity of cutting-plane proofs. *Discret. Appl. Math.*, 18(1): 25–38.
- Cook, W. J.; Koch, T.; Steffy, D. E.; and Wolter, K. 2013. A hybrid branch-and-bound approach for exact rational mixed-integer programming. *Math. Program. Comput.*, 5(3): 305–344.
- Davies, J. 2014. *Solving MAXSAT by Decoupling Optimization and Satisfaction*. Ph.D. thesis, University of Toronto, Canada.
- Davies, J.; and Bacchus, F. 2011. Solving MAXSAT by Solving a Sequence of Simpler SAT Instances. In *CP*, volume 6876 of *LNCS*, 225–239. Springer.
- Davies, J.; and Bacchus, F. 2013a. Exploiting the Power of mip Solvers in MaxSAT. In *SAT*, volume 7962 of *LNCS*, 166–181. Springer.
- Davies, J.; and Bacchus, F. 2013b. Postponing Optimization to Speed Up MAXSAT Solving. In *CP*, volume 8124 of *LNCS*, 247–262. Springer.
- Delisle, E.; and Bacchus, F. 2013. Solving Weighted CSPs by Successive Relaxations. In *CP*, volume 8124 of *LNCS*, 273–281. Springer.
- Devriendt, J.; Gleixner, A. M.; and Nordström, J. 2021. Learn to relax: Integrating 0-1 integer linear programming with pseudo-Boolean conflict-driven search. *Constraints An Int. J.*, 26(1): 26–55.
- Devriendt, J.; Gocht, S.; Demirovic, E.; Nordström, J.; and Stuckey, P. J. 2021. Cutting to the Core of Pseudo-Boolean Optimization: Combining Core-Guided Search with Cutting Planes Reasoning. In *AAAI*, 3750–3758. AAAI Press.
- Eiffler, L.; and Gleixner, A. M. 2021. A Computational Status Update for Exact Rational Mixed Integer Programming. In *IPCO*, volume 12707 of *LNCS*, 163–177. Springer.
- Eiffler, L.; and Gleixner, A. M. 2023. A computational status update for exact rational mixed integer programming. *Math. Program.*, 197(2): 793–812.
- Flippo, M.; Sidorov, K.; Marijnissen, I.; Smits, J.; and Demirovic, E. 2024. A Multi-Stage Proof Logging Framework to Certify the Correctness of CP Solvers. In *CP*, volume 307 of *LIPICs*, 11:1–11:20. Schloss Dagstuhl.
- Gamba, E.; Bogaerts, B.; and Guns, T. 2023. Efficiently Explaining CSPs with Unsatisfiable Subset Optimization. *J. Artif. Intell. Res.*, 78: 709–746.
- Gocht, S.; McCreesh, C.; and Nordström, J. 2022. An Auditable Constraint Programming Solver. In *CP*, volume 235 of *LIPICs*, 25:1–25:18. Schloss Dagstuhl.

- Gocht, S.; and Nordström, J. 2021. Certifying Parity Reasoning Efficiently Using Pseudo-Boolean Proofs. In *AAAI*, 3768–3777. AAAI Press.
- Hyttinen, A.; Saikko, P.; and Järvisalo, M. 2017. A Core-Guided Approach to Learning Optimal Causal Graphs. In *IJCAI*, 645–651. ijcai.org.
- Ignatiev, A.; Morgado, A.; and Marques-Silva, J. 2016. Propositional Abduction with Implicit Hitting Sets. In *ECAI*, volume 285 of *FAIA*, 1327–1335. IOS Press.
- Ignatiev, A.; Previti, A.; Liffiton, M. H.; and Marques-Silva, J. 2015. Smallest MUS Extraction with Minimal Hitting Set Dualization. In *CP*, volume 9255 of *LNCS*, 173–182. Springer.
- Ihalainen, H.; Vandesande, D.; Schidler, A.; Berg, J.; Bogaerts, B.; and Järvisalo, M. 2025. Experimental Repository for “Efficient and Reliable Hitting-Set Computations for the Implicit Hitting Set Approach”. <https://doi.org/10.5281/zenodo.17600095>.
- Jabs, C.; Berg, J.; Bogaerts, B.; and Järvisalo, M. 2025. Certifying Pareto Optimality in Multi-Objective Maximum Satisfiability. In *TACAS (2)*, volume 15697 of *LNCS*, 108–129. Springer.
- Liffiton, M. H.; and Sakallah, K. A. 2008. Algorithms for Computing Minimal Unsatisfiable Subsets of Constraints. *J. Autom. Reason.*, 40(1): 1–33.
- Niskanen, A.; Mustonen, J.; Berg, J.; and Järvisalo, M. 2022. Computing Smallest MUSes of Quantified Boolean Formulas. In *LPNMR*, volume 13416 of *LNCS*, 301–314. Springer.
- Oertel, A. 2025. PBOxide: Verifier for pseudo-Boolean proofs rewritten in Rust. <https://gitlab.com/MIAOresearch/software/pboxide>.
- Roussel, O. 2024. Pseudo-Boolean Competition 2024. <https://www.cril.univ-artois.fr/PB24/>.
- Roussel, O.; and Manquinho, V. M. 2009. Pseudo-Boolean and Cardinality Constraints. In *Handbook of Satisfiability*, volume 185 of *FAIA*, 695–733. IOS Press.
- Saikko, P.; Berg, J.; and Järvisalo, M. 2016. LMHS: A SAT-IP Hybrid MaxSAT Solver. In *SAT*, volume 9710 of *LNCS*, 539–546. Springer.
- Saikko, P.; Dodaro, C.; Alviano, M.; and Järvisalo, M. 2018. A Hybrid Approach to Optimization in Answer Set Programming. In *KR*, 32–41. AAAI Press.
- Saikko, P.; Wallner, J. P.; and Järvisalo, M. 2016. Implicit Hitting Set Algorithms for Reasoning Beyond NP. In *KR*, 104–113. AAAI Press.
- Smirnov, P. 2021. *Pseudo-Boolean Optimization by Implicit Hitting Sets*. Master’s thesis, University of Helsinki.
- Smirnov, P.; Berg, J.; and Järvisalo, M. 2021. Pseudo-Boolean Optimization by Implicit Hitting Sets. In *CP*, volume 210 of *LIPICs*, 51:1–51:20. Schloss Dagstuhl.
- Smirnov, P.; Berg, J.; and Järvisalo, M. 2022. Improvements to the Implicit Hitting Set Approach to Pseudo-Boolean Optimization. In *SAT*, volume 236 of *LIPICs*, 13:1–13:18. Schloss Dagstuhl.
- Vandesande, D.; Coll, J.; and Bogaerts, B. 2026. Certified Branch-and-Bound MaxSAT Solving. In *Proceedings of The 40th Annual AAAI Conference on Artificial Intelligence*. Accepted for publication.
- Vandesande, D.; De Wulf, W.; and Bogaerts, B. 2022. QMaxSATpb: A Certified MaxSAT Solver. In *LPNMR*, volume 13416 of *LNCS*, 429–442. Springer.