# Boosting MCSat Modulo Nonlinear Integer Arithmetic via Local Search

Enrico Lipparini<sup>1,2</sup>, <sup>[⊠]</sup><sup>©</sup>, Thomas Hader<sup>3</sup>, Ahmed Irfan<sup>4</sup><sup>©</sup>, and Stéphane Graham-Lengrand<sup>4</sup><sup>©</sup>

 <sup>1</sup> University of Cagliari, Cagliari, Italy enrico.lipparini@unica.it
 <sup>2</sup> University of Genoa, Genoa, Italy
 <sup>3</sup> TU Wien, Vienna, Austria thomas.hader@tuwien.ac.at
 <sup>4</sup> SRI International, Menlo Park, CA, USA {ahmed.irfan,stephane.graham-lengrand}@sri.com

Abstract. The Model Constructing Satisfiability (MCSat) approach to the SMT problem extends the ideas of CDCL from the SAT level to the theory level. Like SAT, its search is driven by incrementally constructing a model by assigning concrete values to theory variables and performing theory-level reasoning to learn lemmas when conflicts arise. Therefore, the selection of values can significantly impact the search process and the solver's performance. In this work, we propose guiding the MCSat search by utilizing assignment values discovered through local search. First, we present a theory-agnostic framework to seamlessly integrate local search techniques within the MCSat framework. Then, we highlight how to use the framework to design a search procedure for (quantifier-free) Nonlinear Integer Arithmetic ( $\mathcal{NIA}$ ), utilizing accelerated hill-climbing and a new operation called *feasible-sets jumping*. We implement the proposed approach in the MCSat engine of the YICES2 solver, and empirically evaluate its performance over the  $\mathcal{NIA}$  benchmarks of SMT-LIB.

# 1 Introduction

Satisfiability Modulo Theory (SMT) is the problem of deciding the satisfiability of a first-order formula with respect to defined background theories. SMT solvers are the core backbone of a vast range of verification and synthesis tools that require reasoning about expressive logical theories such as real/integer arithmetic [3,16]. One of the major state-of-the-art approaches to SMT is the Model Constructing Satisfiability calculus (MCSat) [29,15], which generalizes the ideas of Conflict-Driven Clause Learning (CDCL) to the theory level, and which has been shown to perform particularly well on complex theories such as nonlinear arithmetic. In the MCSat approach, the solver progressively constructs a *theory model*, similarly to how SAT solvers construct Boolean models. Theory reasoning is used to assess the consistency of partial models, provide explanations of infeasibility, decide theory variables, and propagate theory constraints.

When extending the partial model with a new assignment to a theory variable, picking a good value is critical for the overall performance of the solver. Heuristics used by state-of-the-art solvers pick values on the basis of compatibility with the current search state and of computational cheapness. This has a major drawback: these heuristics only consider knowledge of the current search state, neglecting information on how likely a particular assignment is to lead to a satisfying model eventually.

In this work, we address the problem of choosing good values for variable decisions by augmenting the current search state knowledge with insights provided by local search techniques. Following the logic-to-optimization approach [19,37], we associate to the logical formula a cost function that represents the *distance from a model*, and use local search to find assignments that have a small cost. These assignments are then used to guide future MCSat decisions.

Although local search has already been used in the context of SMT, either as a standalone solver [46,35,6] or as a CDCL(T) theory solver [48], our work is the first to propose a tight integration of local search within the MCSat framework, creating a powerful synergy between the reasoning capabilities of MCSat and the intuition provided by local search which boosts performance for both satisfiable and unsatisfiable instances. Our novel approach is flexible enough to allow calls to local search at any point during the MCSat search, seamlessly fitting with the current state. As MCSat progresses through decisions, propagations, and conflicts, the local search problem is instantiated accordingly: (i) the cost function is built upon the simplification of the original formula under current state assumptions, (ii) initial local search assignments are based on the current search state as well as cached values, and (iii) local search moves are enhanced by information on intervals of feasible assignments to theory variables as tracked by the MCSat engine.

While our approach can be applied to any theory supported by MCSat, in this work we showcase its application to the theory of nonlinear integer arithmetic  $(\mathcal{NIA})$ . In particular, we design a procedure based on a new operation called *feasible-sets jumping*, which allows to move between feasible intervals, and on accelerated hill-climbing [26], to move inside feasible intervals.

Contributions. In this work we: (i) design a theory-agnostic framework to tightly integrate local search techniques within the MCSat approach in order to guide variable decisions, (ii) use the framework to define a local search procedure for the theory of nonlinear integer arithmetic, that makes use of feasible-sets jumping and accelerated hill-climbing, and (iii) show the practical applicability of our method using our implementation in the MCSat engine of YICES2 [17] on the quantifier-free  $\mathcal{NIA}$  benchmark set of SMT-LIB [2].

Structure. In Section 2, we provide the necessary background. Section 3 describes a deep integration of local search techniques within the MCSat framework from a general point of view, which is applied in Section 4 to define a local search approach for non-linear integer arithmetic. In Section 5, we show and discuss the results of our experiments before presenting related work in Section 6 and concluding in Section 7.

# 2 Preliminaries

We assume basic knowledge on the standard first-order quantifier-free logical setting and standard notions of theory, satisfiability, and logical consequence. We write logical variables with  $x, y, \ldots$ , and concrete values with  $\alpha, \beta, \ldots$  (the domain of concrete values is theory specific, e.g.  $\mathbb{Z}$  for integer arithmetic,  $\mathbb{R}$  for real arithmetic). An assignment  $\mu$  is a map from variables to values of matching type. If  $\phi$  is a formula, we denote with  $Vars(\phi)$  the set of its (free) variables. We use Cto denote clauses, and L to denote literals. Nonlinear Integer Arithmetic ( $\mathcal{NIA}$ ) is the theory consisting of arbitrary Boolean combinations of Boolean variables and arithmetic atoms of the form of polynomial equalities and polynomial inequalities over integer variables. The theory is undecidable by Matiyasevich's theorem [40].

## 2.1 SMT & MCSat

SMT [3] is the problem of deciding the satisfiability of a first-order formula with respect to some theory or combination of theories. Two of the major approaches for SMT solving are the Conflict-Driven Clause Learning with theory support (CDCL(T)) [42,3] and the MCSat approach. In the former, theory solvers augment a propositional SAT engine with theory reasoning procedures which are capable of deciding a conjunction of literals (i.e. atomic formulas and their negations) in a particular theory. A propositional model (of the Boolean abstraction of the formula) found by the SAT solver is then checked by all theory engines for theory consistency.

The latter, MCSat, applies CDCL-like mechanisms to perform theory reasoning directly. It can be used either as a theory solver for a specific theory (e.g. in Z3 [14] for non-linear arithmetic over the reals and the integers [4]), or as a fully-fledged stand-alone engine able to handle multiple theories (e.g. in Yices2 for non-linear arithmetic over the reals [30] and over the integers [28], bit-vectors [21], arrays [27], and finite fields [25,24]; as well as in SMT-RAT [13] for non-linear real arithmetic [33]). The MCSat architecture consists of a core solver, an assignment trail, and plugins for theory reasoning. Figure 1 illustrates the high level flow of the MCSat framework.

The core solver incrementally constructs a partial model consisting of Boolean and theory assignments (stored in a *trail*), maintaining the invariant that none of the constraints evaluate to false under the partial model. The trail contains three kinds of elements: *propagated literals* (literals implied to be true by the current state), *decided literals* (literals that we assume to be true), and *model assignments* (assignments of first-order variables to concrete values). Propagations, conflict analysis, lemmas generation, and variable decisions are all handled by theory plugins (including a Boolean plugin that is responsible for propositional reasoning). In general, plugins also keep a *feasibility set* for each variable of their competence, containing the values that are consistent with the current trail and are, thus, candidates to be picked for deciding the variable.



Fig. 1: The MCSat framework consists of the following steps: 1) Propagate the trail. 2) If a conflict is found during propagation, check if there is any decision to backtrack over. If not, return UNSAT. Otherwise, explain the conflict using a lemma, backtrack the trail, and repeat step 1. 3) If no conflict is found during propagation, decide on a variable that is not on the trail. If there is nothing left to decide, return SAT. Otherwise, add the decided variable to the trail and repeat step 1.

When the core solver selects a variable for decision, the choice of the value to assign to the variable is handled by the theory plugin responsible for its type. Some solvers (e.g. Yices2) implement a heuristic called *value cache* (a generalization of SAT phase saving [43]), that keeps track of the last value assigned to a variable when the assignment is undone. Then, when a decision has to be made for the variable, the cached value will be used, provided that it is still in the feasibility set; otherwise, it will simply be ignored, and other heuristics will be used (e.g. picking a default value such as 0).

In the following, we will denote with M the trail, with v[M](x) the value of the variable x in the trail (which may be equal to undef if the variable is not assigned in the trail), and with v[M](L) the value of the literal L under the assignment in M which may be true  $(\top)$  or false  $(\bot)$  if L can be fully evaluated under such assignment, or undef otherwise. We denote with  $feasible_M(x)$  the feasibility set of x in M, i.e. all values that can be chosen for x in the current search state given by M. For arithmetical theories, we have that  $feasible_M(x) \subseteq \mathbb{R}$ , and, in particular, that  $feasible_M(x)$  is the union of a finite set of feasible intervals, i.e.  $feasible_M(x) = \bigcup_{i \in [0:m]} I_i$ . We assume that theory plugins provide a function pick value(S) that returns a value from a set S.

*Example 1.* Assume a search problem in  $\mathbb{Z}$  with variables x, y, and z given by the input formula  $\phi$ .

$$\phi = (\neg (x \ge 1) \lor (xy = 1)) \land (\neg (xy = 1) \lor (x + 2yz > 0)) \land (z^2 > 1)$$

A possible trail at some point during the search is

$$M = [(z^2 > 1) \mapsto \top, x \mapsto 1, (x \ge 1) \mapsto \top, (xy = 1) \mapsto \top]$$

On M elements are either decided or propagated. The feasibility sets are given by  $feasible_M(z) = (-\infty, -1) \cup (1, \infty)$  and  $feasible_M(y) = \{1\} = feasible_M(x)$ . Since  $feasible_M(y)$  is a singleton, we can propagate the assignment  $y \mapsto 1$  on M. We further have that v[M]((x + 2yz > 0)) = undef and v[M](x) = 1.

## 2.2 Local Search

We define a local search problem as a triple  $(\mu_0, f_c, moves)$ , where:

- $-\mu_0$  is an initial assignment for a set of variables Vars,
- $-f_c$  is a cost function from the set of assignments to  $\mathbb{R}_{>0}$ ,
- moves is a neighbor relation between assignments.

A local search algorithm starts from the initial assignment  $\mu_0$  and iteratively explores neighboring assignments according to the moves relation. We say that  $\mu'$  is a move from  $\mu$  if  $(\mu, \mu') \in moves$ . A move  $\mu'$  is accepted if  $f_c(\mu') < f_c(\mu)$ . When a move is accepted, the new assignment becomes the current assignment and the search continues until either: a zero-cost assignment is found, there are no more possible moves (meaning that the current assignment represents a local minimum), or a given stopping criterion is reached (e.g. number of moves).

The problem of finding a solution for an SMT formula  $\phi$  can be encoded as a local search problem, e.g., by following the logic-to-optimization approach [19,37,38], in which a formula  $\phi$  is mapped to a term  $\mathcal{L2O}(\phi)$  that represents the distance from a solution.

In principle, the  $\mathcal{L2O}$  operator can be defined for any theory for which the concept of distance between terms makes sense. Here, we limit ourselves to arithmetic theories. We introduce an arithmetic function symbol d of arity 2 and we assume a fixed interpretation d that satisfies the properties of metric distance, i.e. symmetry, positivity, reflexivity, and triangle inequality. We also assume the existence of a fixed constant term  $\epsilon$ , such that  $\epsilon > 0$ . The specific choice of d and  $\epsilon$  is theory-dependent.

We recursively define  $\mathcal{L2O}$  as follows:

$\mathcal{L}2\mathcal{O}(b)$	$\stackrel{\text{\tiny def}}{=} ITE(b,0,1)$
$\mathcal{L}2\mathcal{O}(\neg b)$	$\stackrel{\text{\tiny def}}{=} ITE(b,1,0)$
$\mathcal{L}2\mathcal{O}(t_1 = t_2)$	$\stackrel{\text{\tiny def}}{=} d(t_1, t_2)$
$\mathcal{L}2\mathcal{O}(t_1 \le t_2)$	$\stackrel{\text{\tiny def}}{=} ITE(t_1 \le t_2, \ 0, \ d(t_1, t_2))$
$\mathcal{L}2\mathcal{O}(t_1 < t_2)$	$\stackrel{\text{\tiny def}}{=} ITE(t_1 < t_2, \ 0, \ d(t_1, t_2) + \epsilon)$
$\mathcal{L}2\mathcal{O}(t_1 \neq t_2)$	$\stackrel{\text{\tiny def}}{=} ITE(t_1 \neq t_2, \ 0, \ 1)$
$\mathcal{L}2\mathcal{O}(\phi_1 \wedge \phi_2)$	$\stackrel{\text{\tiny def}}{=} \mathcal{L}2\mathcal{O}(\phi_1) + \mathcal{L}2\mathcal{O}(\phi_2)$
$\mathcal{L}2\mathcal{O}(\phi_1 \lor \phi_2)$	$\stackrel{\text{\tiny def}}{=} \mathcal{L}2\mathcal{O}(\phi_1) \cdot \mathcal{L}2\mathcal{O}(\phi_2)$
$\mathcal{L}2\mathcal{O}(ITE(\phi_c,\phi_1,\phi_2))$	$\stackrel{\text{\tiny def}}{=} ITE(\phi_c, \ \mathcal{L}2\mathcal{O}(\phi_1), \ \mathcal{L}2\mathcal{O}(\phi_2))$
$\mathcal{L}2\mathcal{O}(\neg ITE(\phi_c, \phi_1, \phi_2))$	$\stackrel{\text{\tiny def}}{=} ITE(\phi_c, \ \mathcal{L}2\mathcal{O}(\neg\phi_1), \ \mathcal{L}2\mathcal{O}(\neg\phi_2))$

It is easy to check that a complete assignment  $\mu$  satisfies  $\phi$  if and only if  $\mathcal{L2O}(\phi)$  evaluates to 0 under  $\mu$ .

In the following, with a slight abuse of notation, we denote with  $\mathcal{L2O}(\phi)$  also the corresponding arithmetic function determined by the interpretation d and the constant  $\epsilon$ , and we define the *cost function* associated to  $\phi$  as  $f_c \stackrel{\text{def}}{=} \mathcal{L2O}(\phi)$ .

*Example 2.* Let  $\phi \stackrel{\text{def}}{=} b \wedge x = y^2$ , and  $d(t_1, t_2) \stackrel{\text{def}}{=} |t_1 - t_2|$ . Then, the cost function associated to  $\phi$  is

$$f_c \stackrel{\text{\tiny def}}{=} \mathcal{L}2\mathcal{O}(\phi) = ITE(b, 0, 1) + |x - y^2|$$

Now, let  $\mu_0 \stackrel{\text{def}}{=} \{b \mapsto \bot ; x \mapsto 4 ; y \mapsto 1\}$  be a starting assignment. We have that  $f_c(\mu_0) = ITE(\bot, 0, 1) + |4 - 1^2| = 1 + 3 = 4$ . If we consider the move  $\mu_1 \stackrel{\text{def}}{=} \{b \mapsto \top ; x \mapsto 4 ; y \mapsto 1\}$  that flips b, then  $f_c(\mu_1) = ITE(\top, 0, 1) + |4 - 1^2| = 0 + 3 = 3$ , hence the move is improving and is accepted. Then, if we consider the move  $\mu_2 \stackrel{\text{def}}{=} \{b \mapsto \top ; x \mapsto 4 ; y \mapsto 2\}$  that increases the value of y by 1, we have  $f_c(\mu_1) = ITE(\top, 0, 1) + |4 - 2^2| = 0 + 0 = 0$ . Hence we have found a zero for  $f_c$ , i.e. a satisfying assignment for  $\phi$ .

In general, local search is not guaranteed to find a solution of  $\phi$ , if there is any. Nevertheless, it returns a *local minimum/best-effort value* of the cost function in the neighborhood of the initial assignment.

## **3** Deep combination of Local Search and MCSat

We propose a deep combination of MCSat and local search where: (i) the current state of MCSat is used to instantiate a local search problem and (ii) the results of the local search help guiding future MCSat decisions. Assuming we have a local search procedure LS, we discuss how to instantiate LS (Section 3.1), as well as how to use the result of LS within MCSat (Section 3.2).

### 3.1 Instantiating the Local Search problem

For the instantiation of LS, we determine the initial assignment and the formula upon which the cost function is constructed. Both choices are of fundamental importance. A good initial assignment is essential to find a good local minimum of the cost function. A *good local minimum* is a local minimum that meets two conditions: (i) it has a smaller cost compared to the cost of the initial assignment and (ii) its assignment values are likely to be accepted by MCSat, i.e., they are consistent with the current trail. Passing a simplified formula that takes the truth value of propagated and decided literals into account is also essential to tailor the search to the current MCSat state and to avoid unnecessary computations.

Initial assignment. For every model assignment  $x \mapsto \alpha$  in M, the assigned variable is treated as a constant that takes its respective assigned value (i.e., x is treated as the constant  $\alpha$ ) and is not allowed to be changed in LS. This reduces the dimension of the LS search space, and avoids moves inconsistent with the

Algorithm 1 Initial assignment for LS

**Input:** a set  $Vars(\phi)$ , a trail M, a value cache cache, a feasibility map  $feasible_M$ **Output:** an initial assignment  $\mu_0$ . a set of fixed variables  $VarsFixed \subseteq Vars(\phi)$ 1:  $VarsFixed \leftarrow \emptyset$ 2: for  $x \in Vars(\phi)$  do if  $v[M](x) \neq \text{undef then}$  $\triangleright$  check if x has a value in the trail 3: 4:  $\mu_0(x) \leftarrow v[M](x)$ ▷ assign trail value VarsFixed.add(x)5:else if  $cache(x) \neq undef$  and  $cache(x) \in feasible_M(x)$  then 6: 7:  $\mu_0(x) \leftarrow cache(x)$  $\triangleright$  assign feasible cached value 8: else 9:  $\mu_0(x) \leftarrow pick\_value(feasible_M(x))$  $\triangleright$  assign any feasible value

current trail. For initial assignment of variables that are unassigned in M, a reasonable choice is to use cached values of previous search states, if present in the value cache. However, cached values are not guaranteed to be in the feasibility set, as they might be the result of a previous decision that eventually led to a conflict. Hence, we first check if the cached value is feasible. If it is not, or there is no cached value, we pick any value from the feasibility set by asking the appropriate theory plugin. The procedure for choosing the initial assignment is shown in Algorithm 1. Note that, for all the variables, the feasibility set cannot be empty. An empty feasibility set indicates an inconsistent trail which is resolved using conflict resolution before starting LS. MCSat maintains the invariant that, for a consistent trail, the feasibility set of all variables is non-empty.

Formula for LS. Every Boolean assignment  $L \mapsto \{\top, \bot\}$  in M represents the truth value of the literal L that is assumed to hold at the current search state

Algorithm 2 Formula for LS	
<b>Input:</b> a formula $\phi$ , a trail M	
<b>Output:</b> a subformula $\phi_{LS}$ of $\phi$	
1: $\phi_{LS} \leftarrow \top$	$\triangleright$ formula to be passed to LS
2: for $C \in \phi$ do	
3: $C_{LS} \leftarrow \bot$	
4: for $L \in C$ do	
5: <b>if</b> $v[M](L) = \top$ <b>then</b>	$\triangleright$ if literal is assigned to true in trail
6: $C_{LS} \leftarrow \top$	$\triangleright$ substitute the clause with true
$7: \qquad \phi_{LS} \leftarrow \phi_{LS} \wedge L$	$\triangleright$ store literal
8: break	
9: else if $v[M](L) = \bot$ then	$\triangleright$ if literal is assigned to false in trail
10: $\phi_{LS} \leftarrow \phi_{LS} \land \neg L$	$\triangleright$ store literal (with correct polarity)
11: continue	$\triangleright$ ignore literal in the clause
12: <b>else</b>	
13: $C_{LS} \leftarrow C_{LS} \lor L$	$\triangleright$ keep literal in the clause
14: $\phi_{LS} \leftarrow \phi_{LS} \wedge C_{LS}$	$\triangleright$ store simplified clause

(either because of a propagation or a decision). We can use this information to simplify the original formula before passing it to LS. For a given clause  $C \stackrel{\text{def}}{=} L \vee L_1 \vee \ldots$ , if  $v[M](L) = \top$ , then, for LS, it suffices to find an assignment that satisfies L, since such an assignment would satisfy C as well. Hence, in this case, we shall pass to LS just L instead of C. On the other hand, if  $v[M](L) = \bot$ , then there is no incentive for LS to try to find an assignment that makes L true, as any such assignment would be inconsistent with the trail and will be discarded by MCSat immediately. Thus, L is removed from the clause that is passed to LS. Note that, by just removing L from the clause, we still may end up with an assignment that evaluates L to  $\top$ . Therefore, for literals that are assigned to  $\bot$  in the trail, we add, just once,  $\neg L$  to the formula that we pass to LS. This procedure is shown in Algorithm 2.

## 3.2 Guiding MCSat decisions

During the search, we periodically call LS to suggest values for MCSat to choose in subsequent decisions. As a heuristic to decide when to call LS, we are utilizing a polynomially increasing conflict threshold. Initially, this limit is set to 50, and then it is increased according to the polynomial  $50 \cdot ls_calls \cdot log_{10}(ls_calls + 9)^3$ , where  $ls_calls$  represents the number of times LS has been called. A similar heuristic is used by SAT solvers to decide when to perform certain cache clearing operations [8]. Once the threshold is reached, we wait until the last conflict has been resolved and all consequences of that conflict are propagated. Then we start LS to guide any further decisions.

The return of LS consists of a complete assignment that contains *suggested* values for future variable decisions. These suggested values are put in the MCSat value cache – recall that the values in the cache are picked first during variable decisions, provided that they are feasible. Note that, during the choice of the initial assignment to pass to LS, we had relied on (feasible) cached values as well. If a cached value was feasible and LS changed its value, it means that such change led to a smaller cost, hence got us closer to a solution. Therefore, replacing the old cached value with the newly found suggestion improves the cache quality. On the other hand, if the cached value was not feasible, then any change to a feasible value improves the cache quality.

Furthermore, LS keeps track of the activity of each variable during its execution. The most active variables have contributed most to the decrease of the cost during LS. We suggest those variables to MCSat as good choices for subsequent decisions. This way, variables that were more active during the local search phase will have a higher impact on the MCSat search.

# 4 Local Search for Nonlinear Integer Arithmetic

As explained in Section 3.1, LS receives an initial assignment  $\mu_0$  and a formula  $\phi$  from MCSat. The formula  $\phi$  is used to construct the cost function  $f_c$  using the logic-to-optimization approach (Section 2.2), i.e.  $f_c \stackrel{\text{def}}{=} \mathcal{L2O}(\phi)$ . To apply that,

we must first define the distance function d and the strict inequality constant  $\epsilon$  for integers. For d, we choose a consistent and computationally cheap definition  $d(t_1, t_2) \stackrel{\text{def}}{=} |t_1 - t_2|$ . For  $\epsilon$ , our choice is  $\epsilon \stackrel{\text{def}}{=} 1$ , since t < 0 is interchangeable with  $t + 1 \leq 0$  for integers.

The building blocks of local search are *moves*. We contemplate three types of moves (or *modes*): one for Boolean variables, and two for integer variables. Given an assignment  $\mu$ , we have the following types of moves:

- Boolean flips: For a Boolean variable b, the assignment  $\mu_{\neg b} \stackrel{\text{def}}{=} \mu[b \mapsto \neg \mu(b)]$ obtained by changing the value of b to the negation of its value assigned by  $\mu$  is a flip move from  $\mu$ .
- Hill-climbing moves: In the basic version of hill-climbing, for an integer variable x, the assignments  $\mu_{x+1} \stackrel{\text{def}}{=} \mu[x \mapsto \mu(x) + 1]$  and  $\mu_{x-1} \stackrel{\text{def}}{=} \mu[x \mapsto \mu(x) 1]$  obtained by mapping x to the successor and predecessor of its value assigned by  $\mu$  are moves from  $\mu$ .
- Feasible-set-jumps: For an integer variable x, with feasibility set  $feasible(x) = \bigcup_{i \in [0:m]} I_i$ , and  $x \in I_j$  (for a given  $j \in [0:m]$ ), the assignments  $\mu_{\text{left}} \stackrel{\text{def}}{=} \mu[x \mapsto pick\_value(I_{j-1})]$ , and  $\mu_{\text{right}} \stackrel{\text{def}}{=} \mu[x \mapsto pick\_value(I_{j+1})]$  obtained by picking a value from the left and right feasible intervals of  $I_j$  (provided they exist, i.e., respectively, that  $j-1 \in [0:m]$ , and  $j+1 \in [0:m]$ ) are moves from  $\mu$ .

Our specific strategy of the local search algorithm is outlined in Algorithm 3. The algorithm starts with a list of variables vars (and an associated *feasible* 

## Algorithm 3 LS main algorithm

Inp	put: a list vars, a feasibility map feasible	e, an initial assign. $\mu_0$ , a cost function $f_c$
Ou	<b>itput:</b> a final assignment $\mu^*$ with $f_c(\mu^*)$	$\leq f_c(\mu_0)$
1:	$\mu^* \leftarrow \mu_0$	$\triangleright$ best assignment
2:	$cost^* \leftarrow f_c(\mu_0)$	$\triangleright$ best cost
3:	for $mode \in \{bool-flips, fs-jumps, hill-clips, hill-c$	nb} do
4:	$n\_vars \leftarrow 0$ >	no. of vars visited since last improvement
5:	while $n$ vars $< len(vars)$ and cost	$* \neq 0$ do
6:	$x \leftarrow vars[n\_vars]$	$\triangleright$ pick next variable
7:	$lpha \leftarrow \mu^*(x)$	$\triangleright$ value assigned to x
8:	while $\alpha_{new} \leftarrow \text{Move.choose}(x, \alpha)$	a, feasible, mode) do
9:	$\mu_{new} \leftarrow \mu^*[x \mapsto \alpha_{new}]$	$\triangleright$ create new assignment
10:	$cost_{new} \leftarrow f_c(\mu_{new})$	
11:	$success \leftarrow cost_{new} < cost^*$	$\triangleright$ check if the move has improved
12:	if success then	
13:	$\mu^* \leftarrow \mu_{new}$	$\triangleright$ update best assignment
14:	$cost^* \leftarrow cost_{new}$	$\triangleright$ update best cost
15:	$n\_vars \leftarrow 0$	$\triangleright$ reset no. of vars visited
16:	$vars.to\_front(x)$	$\triangleright$ move x to the front of the list
17:	MOVE.notify $(x, \alpha, \alpha_{new}, feasily)$	ble, mode, success)
18:	$n\_vars \leftarrow n\_vars + 1$	$\triangleright$ increas no. of vars visited

map), an initial assignment  $\mu_0$ , and a cost function  $f_c$ . The goal of the procedure is to return an assignment  $\mu^*$  that improves over the initial assignment  $\mu_0$  w.r.t. the cost function, i.e.  $f_c(\mu^*) < f_c(\mu_0)$ .

At the beginning, the best assignment coincides with the initial assignment (Line 1). First, we cycle over modes (Line 3). Then, we enter in a loop over the variables (Line 5). The loop breaks only in two cases: if all the variables have already been visited since the last improvement (in which case, it means we have reached a local minimum w.r.t. the current mode moves), or if the current cost is equal to 0 (in which case it means we have found a solution). At each loop iteration, we pick the next variable (Line 6). Here, for simplicity, we assume that there are no fixed variables (in practice, these variables are just ignored and treated as constants). Then, for the current variable x, we enter in a second loop (Line 8), in which we select new values for x. These values are determined by a move selection module, which we discuss below. The loop breaks only when there are no more moves available. For each value, a new assignment is built by re-assigning x to the new value (Line 9), and the cost of the new assignment is computed (Line 10). We then check whether the current cost is lower than the previous cost (Line 11). If so, then the new assignment becomes the best assignment (Line 13), and x is moved to the front of the list (Line 16). If not, then we try other moves for x, if there are any. In both cases, we notify the move selection module whether the suggested move has led to a success or not (Line 17). The move selection module works as following.

*Boolean flips mode.* Here, the logic is rather straightforward, as there is only one move possible per variable. Regardless of whether the move has success or not, the cycle over moves terminates, and the algorithm proceeds with the cycles over variables or over modes.

Accelerated hill-climbing mode. The simple hill-climbing moves presented earlier, in which we add or subtract 1 to the current value, can be quite slow in converging toward a local minimum when the search space is huge. For this reason, we accelerate hill-climbing by keeping, for each variable, an adaptive step\_size, which is incremented or decremented according to a fixed acceleration parameter acc (in our setting acc = 1.2) and on the base of the success of previous moves. At the beginning,  $step\_size$  is set to 1 (i.e., we start with simple hill-climbing moves). At each iteration, we try four moves, corresponding to adding to the current value the product between  $step\_size$  and one of the following:  $acc, \frac{1}{acc}, \frac{-1}{acc}, -acc$ . Since we are working with integers, every step value is rounded to the nearest integer. If one of the moves has success, then we set  $step\_size$  to be equal to the best successful step (thus keeping the best velocity). If none of the moves has success, then we stop the moves cycle, and we set  $step\_size$  to  $\frac{step\_size}{acc}$  (thus decelerating over this variable for future moves).

*Feasible-set-jumping mode.* There are two possible versions of fs-jumping: *global* and *local.* In global fs-jumping, given a fixed variable, we try *all* possible jumps over the feasibility set, i.e. we try one jump per feasible interval. While this

11

may give a wide-ranging view over the feasibility set, it can also be very costly, hence we limit global fs-jumping to one time per variable (per LS call). Local fs-jumping, on the contrary, only explores the left and the right feasible intervals w.r.t. to the interval that contains the current value. If one fs-jump is successful, e.g. the one to the left interval, then we continue on that direction and explore the interval further left. As soon as we find that both left and right fs-jumps do not improve, then we stop, hence avoiding to span over all feasible intervals like in the global fs-jumping.

# 5 Experiments

Implementation. We have implemented our method in the MCSat engine of the YICES2 SMT solver, adding a module for the interaction with LS. We will denote the version of YICES2 that makes use of LS as YICES2<sub>LS</sub> and the baseline version (without any local search) as YICES2<sub>base</sub>.

Setup. We have run our experiments on a cluster equipped with AMD EPYC 7502 CPUs running at 2.5GHz, using a timeout of 300 seconds, and a memory limit of 8GB. We have compared the base version of YICES2 with the LS-boosted version YICES2<sub>LS</sub> as well as with the state-of-the-art SMT solvers CVC5 [1] (version 1.2.0), MATHSAT5 [11] (version 5.6.11), and Z3 [14] (version 4.13.3). We have also included in the comparison HYBRIDSMT [7], which runs a portfolio composed by the LS solver LocalSMT [6] (used as a standalone tool) and a LocalSMT-boosted version of Z3's CDCL(T) – see also related work (Section 6).

Benchmarks. We have considered all the SMT-LIB [2] (Version 2024 [44]) benchmarks from the QF\_NIA category. This is a class of 25443 benchmarks, among which 14990 and 5183 come with a known status of "sat" and "unsat", respectively, and another 5270 have an "unknown" satisfiability status.

Results. In the presentation of the results, we consider both a short time limit and a long time limit. We set the short time limit to 24s, as in the respective SMT-COMP track [47], and the long time limit to 300s, due to resource constraints. The results are shown in Table 1 and Table 2, respectively. In the columns, we separate per benchmark family; on the rows, for each solver, we report the amount of overall benchmarks solved, and, in parenthesis, the amount of benchmarks solved restricted to sat and unsat instances, respectively. We also include two portfolios between Z3 (resp., HYBRIDSMT) and YICES2<sub>LS</sub>, that work as follows: we run Z3 (resp., HYBRIDSMT) for half of the time limit (i.e., 12s/150s), then, if it has not terminated, we run YICES2<sub>LS</sub> for the remaining time.

Discussion. First, we observe that, with both time limits,  $Y_{ICES}2_{LS}$  solves a significant number of benchmarks more than  $Y_{ICES}2_{base}$ . In particular, on both satisfiable and unsatisfiable instances, it improves (or matches)  $Y_{ICES}2_{base}$  results over all families, except one. Improving on unsatisfiable benchmarks is

noteworthy: indeed, while, in general, local search is geared toward proving satisfiability, integrating it within MCSat enables to generate better lemmas. This is witnessed not only by the higher amount of benchmarks solved overall, but also by the lower amount of conflicts and theory variables assignments occurred. On unsatisfiable benchmarks solved by both tools, on average (resp. median),

	10 6	Hybered Ship	Marth Star	T Contraction of the second	Land	ŝ	000 × 0000 × 000 × 000 × 000 × 000 × 000 × 000 × 000 × 000 ×	X Property Control of
VeryMax	$7799 \\ (5465) \\ (2334)$	$\begin{array}{c} 13555 \\ (\underline{10427}) \\ (3128) \end{array}$	$\begin{array}{c} 11233 \\ (7615) \\ (3618) \end{array}$	$\begin{array}{c} 13695 \\ (9461) \\ (4234) \end{array}$	$\begin{array}{c} \underline{14269} \\ (10019) \\ (4250) \end{array}$	$\begin{array}{c} 13975 \\ (9599) \\ (\underline{4376}) \end{array}$	$\begin{array}{r} 14848 \\ (10271) \\ (4577) \end{array}$	$\begin{array}{r} 15428 \\ (11105) \\ (4323) \end{array}$
calypto	$171 \\ (79) \\ (92)$	$     \begin{array}{r}       174 \\       (78) \\       (\underline{96})     \end{array} $	$168 \\ (79) \\ (89)$	$174 \\ (79) \\ (95)$	$174 \\ (79) \\ (95)$	$\frac{176}{(\underline{80})}$ ( <u>96</u> )	$     \begin{array}{r}       176 \\       (80) \\       (96)     \end{array} $	$     \begin{array}{r}       176 \\       (80) \\       (96)     \end{array} $
$\mathbf{ezsmt}$	$(\underline{\underline{8}})$ $(\underline{0})$	( <u>8</u> ) ( <u>0</u> )	$(\underline{\underline{8}})$ $(\underline{0})$	( <u>8</u> ) ( <u>0</u> )	( <u>8</u> ) ( <u>0</u> )	8 ( <u>8</u> ) ( <u>0</u> )	8 (8) (0)	8 (8) (0)
LassoRank	97 $(\underline{4})$ (93)	$\begin{array}{c} \underline{105} \\ (\underline{4}) \\ (\underline{101}) \end{array}$	$\frac{105}{(\underline{4})}$ $(\underline{101})$	$91 \\ (\underline{4}) \\ (87)$	$96 \\ (\underline{4}) \\ (92)$	$\begin{array}{c} 104 \\ (\underline{4}) \\ (100) \end{array}$	$     \begin{array}{c}       102 \\       (4) \\       (98)     \end{array} $	$     \begin{array}{r}       104 \\       (4) \\       (100)     \end{array} $
Dartagnan	$320 \\ (11) \\ (309)$	$     \frac{354}{(10)}     (344) $	$327 \\ (\underline{12}) \\ (315)$	$ \begin{array}{c} 142 \\ (1) \\ (141) \end{array} $		$352 \\ (9) \\ (343)$	$ \begin{array}{r}     344 \\     (9) \\     (335) \end{array} $	347 (9) (338)
LCTES	$\begin{array}{c} \underline{1} \\ (\underline{0}) \\ (\underline{1}) \end{array}$	$\begin{array}{c} \underline{1} \\ (\underline{0}) \\ (\underline{1}) \end{array}$	$\begin{array}{c} \underline{1} \\ (\underline{0}) \\ (\underline{1}) \end{array}$	$\begin{array}{c} 0 \\ (\underline{0}) \\ (0) \end{array}$	$\begin{array}{c} 0 \\ (\underline{0}) \\ (0) \end{array}$	$\begin{array}{c} \underline{1} \\ (\underline{0}) \\ (\underline{1}) \end{array}$	$\begin{pmatrix} 1 \\ (0) \\ (1) \end{pmatrix}$	$ \begin{array}{c} 1 \\ (0) \\ (1) \end{array} $
MathProbl	$107 \\ (100) \\ (7)$	$93 \\ (86) \\ (7)$	$ \begin{array}{r}     141 \\     (134) \\     (7) \end{array} $	$     \begin{array}{r}       122 \\       (115) \\       (7)     \end{array} $	$\frac{303}{(296)}$ (7)	$     \begin{array}{r}       118 \\       (111) \\       (7)     \end{array} $	300 (293) (7)	305 (298) (7)
leipgiz	(70) (2)	$(151 \\ (150) \\ (1)$	(114) (112) (2)	$102 \\ (101) \\ (1)$	(110) (1)	$     \begin{array}{r}       120 \\       (119) \\       (1)     \end{array} $	119 (118) (1)	$     \begin{array}{r}       148 \\       (147) \\       (1)     \end{array} $
UltAut23	$(\underline{8})$ (8)	$ \begin{array}{c} 10 \\ (7) \\ (3) \end{array} $	(10)	7 (7) (0)	7 (7) (0)	$\frac{\underline{21}}{(\underline{8})}$ (13)	$ \begin{array}{c} 20 \\ (8) \\ (12) \end{array} $	9 (7) (2)
mcm	9 (9) (0)		3 (3) (0)	6 (6) (0)	6 (6) (0)	5 (5) (0)	5 (5) (0)	46 (46) (0)
$\operatorname{sqrtmodinv}$	$2 \\ (0) \\ (2)$	$\begin{array}{c} 7\\(\underline{0})\\(7)\end{array}$	$\begin{array}{c} 0\\ (\underline{0})\\ (0) \end{array}$	$\begin{array}{c} 0 \\ (\underline{0}) \\ (0) \end{array}$	$\begin{array}{c} 0 \\ (\underline{0}) \\ (0) \end{array}$	$\begin{array}{c} \underline{17} \\ (\underline{0}) \\ (\underline{17}) \end{array}$		6 (0) (6)
UltAut	$(\underline{0})$ $(\underline{7})$	$\begin{array}{c} \underline{7} \\ (\underline{0}) \\ (\underline{7}) \end{array}$	$(\underline{0})$ $(\underline{7})$	$\begin{array}{c} \frac{7}{(\underline{0})}\\ (\underline{7}) \end{array}$	$(\underline{0})$ $(\underline{7})$	$\begin{array}{c} \underline{7} \\ (\underline{0}) \\ (\underline{7}) \end{array}$	7 (0) (7)	7 (0) (7)
AProVE	$\begin{array}{r} 1816 \\ (1251) \\ (565) \end{array}$	$\begin{array}{c} 2212 \\ (1572) \\ (640) \end{array}$	$\begin{array}{r} 2085 \\ (1556) \\ (529) \end{array}$	$2328 \\ (1615) \\ (713)$	$\begin{array}{c} {\color{red} {\bf 2356}} \\ {\color{red} {(\underline{1642})}} \\ {\color{red} {(\underline{714})}} \end{array}$	$\begin{array}{r} 2289 \\ (1612) \\ (677) \end{array}$	$2380 \\ (1657) \\ (723)$	$2371 \\ (1649) \\ (722)$
UltLasso	$\begin{array}{c} \underline{32} \\ (\underline{6}) \\ (\underline{26}) \end{array}$	$\begin{array}{c} 31 \\ (\underline{6}) \\ (25) \end{array}$	$\begin{array}{c} \underline{32} \\ (\underline{6}) \\ (\underline{26}) \end{array}$	$\begin{array}{c} \underline{32}\\ (\underline{6})\\ (\underline{26}) \end{array}$	$\begin{array}{c} \underline{32} \\ (\underline{6}) \\ (\underline{26}) \end{array}$	$\begin{array}{c} \underline{32}\\ (\underline{6})\\ (\underline{26}) \end{array}$	$ \begin{array}{c} 32 \\ (6) \\ (26) \end{array} $	32 (6) (26)
Total	$\begin{array}{r} 10457 \\ (7011) \\ (3446) \end{array}$	$\begin{array}{c} 16764 \\ (\underline{12404}) \\ (4360) \end{array}$	$\begin{array}{r}14241\\(9536)\\(4705)\end{array}$	$\begin{array}{c} 16714 \\ (11403) \\ (5311) \end{array}$	$\begin{array}{c} {\color{red} \frac{{\color{black} {\bf 17456}}}{(12177)}}\\ {\color{black} (5279)}\end{array}$	$\begin{array}{c} 17225 \\ (11561) \\ (\underline{5664}) \end{array}$	$\begin{array}{c} 18359 \\ (12459) \\ (5900) \end{array}$	$\begin{array}{r} 18988 \\ (13359) \\ (5629) \end{array}$

Table 1: Summary of results for  $\mathcal{NIA}$  benchmarks with a timeout of 24s.

13

YICES2<sub>LS</sub> encountered 225 (resp. 4) fewer conflicts and 27826 (resp. 36) fewer theory variables assignments than YICES2<sub>base</sub>. Note that, on average (resp. median), YICES2<sub>base</sub> encountered 1669 (resp. 458) conflicts and 79842 (resp. 10537) theory variable assignments. These numbers show that there is a considerable amount of benchmarks for which the number of conflicts and theory variable

		M.T.	Ŕ	ç	ó		20 -	2417
	6	in the	P.S.A.	1200	AL.		A A AL	Contraction of the
	ê G	47.	$M_{4}$	7. <sup>07</sup>	70	ŵ	2000	EX.
New Me	10489	$\frac{17060}{(10002)}$	13651	14549	15160	16304	16698	17303
verymax	(7122) (3367)	$(\frac{12095}{4967})$	(9309) (4142)	(4385)	(4441)	(11044) (5260)	(5332)	(5147)
	173	175	169	174	175	177	177	177
calypto	(79) (04)	(78)	(79)	(79)	(79)	$(\frac{80}{97})$	(80)	(80)
	8	( <u>91</u> ) 8	(90)	(95)	(90)	( <u>97</u> ) 8	(97)	(97)
ezsmt	$(\overline{\underline{8}})$	$(\overline{\underline{8}})$	$(\overline{\underline{8}})$	$(\overline{8})$	$(\overline{8})$	$(\overline{\underline{8}})$	(8)	(8)
	$(\underline{0})$	$(\underline{0})$	$(\underline{0})$	$(\underline{0})$	$(\underline{0})$	$(\underline{0})$	(0)	(0)
	98	106	105	93	97	106	104	105
LassoRank	$(\underline{4})$	$(\underline{4})$	$(\underline{4})$	$(\underline{4})$	$(\underline{4})$	$(\underline{4})$	(4)	(4)
	(94)	( <u>102</u> ) <b>260</b>	(101)	(89)	(93)	(102)	(100)	(101)
Dartagnan	(17)	(14)	(18)	(7)	(3)	(14)	(13)	(13)
Durtughun	(333)	$(\underline{355})$	(329)	(304)	(285)	(354)	(350)	(354)
	1	2	1	0	0	2	1	1
LCTES	$(\underline{0})$	$(\underline{0})$	$(\underline{0})$	$(\underline{0})$	$(\underline{0})$	$(\underline{0})$	(0)	(0)
	(1)	( <u>2</u> )	(1)	(0)	(0)	( <u>2</u> )	(1)	(1)
MathProbl	(170)	(100)	(176)	124 (117)	$(\frac{311}{304})$	(111)	314 (307)	32( (320)
Mathr 1001	(170)	(100) (7)	(170) (7)	(117)	$(\frac{304}{(7)})$	(111) (7)	(307)	(320)
	89	156	126	104	111	135	134	154
leipgiz	(87)	$(\overline{155})$	(124)	(103)	(110)	(134)	(133)	(153)
	$(\underline{2})$	(1)	$(\underline{2})$	(1)	$(1)^{-1}$	(1)	(1)	(1)
	16	10	17	7	7	21	21	10
UltAut23	( <u>8</u> )	(7)	(7)	(7)	(7)	( <u>8</u> )	(8)	(7)
	(0)	(3)	(10)	(0)	(0)	( <u>13</u> )	(13)	(3)
mam	(17)	(60)	10 (10)	(0)	(0)	10 (10)	(11)	64 (64)
mem	(11) (0)	$\left( \begin{array}{c} (\underline{0} \underline{3}) \\ (\underline{0}) \end{array} \right)$	(0)	$\begin{pmatrix} 0 \\ 0 \end{pmatrix}$	(0)	(10)	(0)	(04)
	2	10	0	0	0	17	17	8
sqrt mod inv	$(\underline{0})$	( <u>0</u> )	$(\underline{0})$	( <u>0</u> )	( <u>0</u> )	$\overline{(\underline{0})}$	(0)	(0)
	(2)	(10)	(0)	(0)	(0)	( <u>17</u> )	(17)	(8)
	<u>7</u>	$\frac{7}{2}$	<u>7</u>	<u>7</u>	<u>7</u>	<u>7</u>	7	7
UltAut	$\left(\frac{0}{7}\right)$	$\left( \begin{array}{c} (\underline{0}) \\ (\underline{7}) \end{array} \right)$	$\left(\frac{0}{7}\right)$	$\left(\begin{array}{c} \underline{0}\\ (\underline{7}) \end{array}\right)$	$\left(\frac{0}{7}\right)$	$(\underline{0})$	(0)	(0)
	<u>(1)</u> 1939	( <u>1</u> ) 2352	<u>(1)</u> 2180	( <u>1</u> ) 2337	2367	<u>(1)</u> 2330	2394	2387
AProVE	(1330)	(1640)	(1622)	(1621)	(1647)	(1640)	(1661)	(1656)
	(609)	(685)	(558)	(716)	(720)	(699)	(733)	(731)
	32	31	32	32	32	32	32	32
$\mathrm{UltL}\mathrm{asso}$	$(\underline{6})$	$(\underline{6})$	$(\underline{6})$	( <u>6</u> )	( <u>6</u> )	( <u>6</u> )	(6)	(6)
	( <u>26</u> )	(25)	( <u>26</u> )	( <u>26</u> )	( <u>26</u> )	( <u>26</u> )	(26)	(26)
	13398	20435	16836	17755	18572	19644	20281	20950
Total	(8848)	$\left  \left( \frac{14174}{(2224)} \right) \right $	(11563)	(12125)	(12896)	(13059)	(13597)	(14467)
	(4550)	(6261)	(5273)	(5630)	(5676)	(6585)	(6684)	(6483)

Table 2: Summary of results for  $\mathcal{NIA}$  benchmarks with a timeout of 300s.

assignments is significantly lower (note that such a lower median w.r.t. average implies a pronounced right-skewness).

Overall, we see that, in the 24s track,  $YICES2_{LS}$  solves more benchmarks than any other solver, while, in the 300s track, it comes third, after HYBRIDSMT and Z3. The complementarity of  $YICES2_{LS}$  w.r.t. both tools can be witnessed by the scatter plots in Figure 3, and by the results of the portfolios in Tables 1 and 2. Note that Z3 internally utilizes portfolio tactics that combine multiple solving techniques sequentially (clearly observable in Figures 2 and 3). HYBRIDSMT runs a higher-level portfolio that combines the LS-based LocalSMT and Z3. The results of the portfolios that include  $YICES2_{LS}$  show that our approach brings significant diversity to the strategies already used in state-of-the-art portfolio approaches.

Since HYBRIDSMT is the only other solver that – to the best of our knowledge – leverages local search techniques for  $\mathcal{NIA}$ , it is interesting to compare the improvements it brings to Z3 with the improvements that  $Y_{ICES2_{LS}}$  brings over  $Y_{ICES2_{base}}$ . We can see that, with a 300s time limit, the improvements are comparable, as both tools solve around 800 benchmarks more than their base solvers. With a 24s time limit, however, we see that  $Y_{ICES2_{LS}}$  is able to solve around 700 benchmarks more than  $Y_{ICES2_{base}}$ , while, on the contrary, HYBRIDSMT loses around 450 benchmarks compared to Z3. Figure 2 shows that the point at which using local search pays off is much earlier for  $Y_{ICES2_{LS}}$  (< 10s) than for HYBRIDSMT (just below 100s).



Fig. 2: Plots showing the number of instances solved (x axis) within given time in seconds (y axis) in log scale.



Fig. 3: Scatter plots comparing YICES  $2_{LS}$  to CVC5, HYBRIDSMT, MATHSAT5, YICES  $2_{base}$ , and Z3, respectively; on sat. (orange) and unsat. (blue) instances.

# 6 Related work

In propositional SAT solving, local search techniques have been successfully used to solve difficult satisfiable problems [31] as well as unsatisfiable instances [45]. Recently, their tight integration in the propositional CDCL framework has been shown to improve performance [7,8] and are now considered a key component of state-of-the-art SAT solvers. In the context of SMT, on the other hand, the adoption of LS is a lot less widespread.

In [22], the LS-based SAT solver WalkSAT has been used in combination with a theory solver as an alternative to the classic CDCL(T) approach; however, the use of local search remained limited to the Boolean level. For the theory of bit-vectors, the idea of Boolean flips in SAT solving has been transposed to the bit level by introducing *bit-flips* [18], possibly augmented with propagations [41].

The adoption of LS for arithmetic theories is more recent. For the theories of Linear Integer Arithmetic ( $\mathcal{LIA}$ ) [5] and Multi-linear Real Arithmetic [34] a critical move operation is used to change the value of a variable that appears in a literal violated by the current assignment in order to make the literal satisfied. To deal with the nonlinear arithmetic constraints, the cell-jumping technique is used, which first isolates the roots of a falsified polynomial w.r.t. to a variable (by fixing the value of the other variables), thus decomposing the real space into finitely many intervals (cells, in the CAD [12] terminology), and then tries to satisfy the polynomial by changing the value of the variable by jumping around these cells. This technique has been implemented in LocalSMT for  $\mathcal{NIA}$  [6], and as a tool in Maple [35] and on top of Z3 [46] for  $\mathcal{NRA}$ .

Local search has also been used as a sub-routine for global search techniques, as in the case of floating points [19], and of  $\mathcal{NRA}$  possibly augmented with transcendental functions ( $\mathcal{NTA}$ ) [37,39,36]. In these works, numerical optimization algorithms, e.g. the gradient-descent, are used to find local minima, while stochastic jumping is used to move away from a local minimum in order to explore other regions in search for a global minimum.

All the methods discussed so far for arithmetic theories are only able to prove satisfiability; if they fail, then all the knowledge that has been acquired by the search is lost. HybridSMT [48] addresses this issue, for the case of  $\mathcal{NIA}$ , by integrating LocalSMT within Z3's CDCL(T). In particular, LocalSMT takes as input a subformula corresponding to a Boolean skeleton solution, and, if it does not find an integer solution for the subformula, it returns the best assignment found and the conflict frequency for atoms. This information is used to improve phase selection (i.e. Boolean assignments) and variable ordering. Although both HybridSMT and our method share the idea of integrating LS within a reasoning calculus (CDCL(T) and MCSat, respectively), there are some substantial difference. First, in HybridSMT, LS takes into account complete Boolean variable assignments. In our framework, LS can take as input both Boolean and theory variable assignments, either partial or complete. Additionally, while in HybridSMT LS can only suggest assignments for (and ordering of) Boolean literals, we extend that to theory variables as well. Moreover, there is a theoryspecific difference in our approach. LocalSMT relies on cell-jumps, which require

17

to perform potentially very expensive root isolation sub-routines at every step. In contrast, our method uses fs-jumps that rely on feasibility intervals already maintained by the theory plugin in the MCSat framework. This eliminates the need for additional computation and can be viewed as a lazy version of cells, progressively refined on-demand. Furthermore, we pair fs-jumps with hill-climbing to move inside feasible intervals.

Most state-of-the-art solvers do not use local search for  $\mathcal{NIA}$  problems. Bitblasting [20] aims at proving satisfiability by iteratively imposing bounds on the variables and then encoding the obtained sub-formula into an equi-satisfiable Boolean formula, which is then handled by a SAT solver. In the branch-andbound approach [32,28] the integer domain is relaxed by allowing variables to range over real numbers. Incremental Linearization [9,10] leverages decision procedures for  $\mathcal{LIA}$  by abstracting non-linear multiplications with uninterpreted functions and then incrementally axiomatize them.

# 7 Conclusion

In this work, we have introduced a theory-independent framework for integrating local search into the MCSat calculus. By combining local search intuition with MCSat reasoning capabilities, our approach leverages logic-to-optimization formalization to provide guidance to the core MCSat solver. Specifically, we addressed the theory of nonlinear integer arithmetic by proposing a local search procedure based on feasibility-set jumping and hill-climbing.

We implemented our approach in the YICES2 SMT solver, and empirically demonstrated its improvements for both satisfiable and unsatisfiable instances. Our results show that the new YICES2 solver with local search compares favorably and often outperforms other SMT solvers; in particular, it manages to solve a significant amount of benchmarks not solved by other state-of-the-art tools. Moreover, our results show that the new YICES2 is able to reach solutions or proofs more efficiently, compared to the baseline YICES2, in terms of the number of decisions and conflicts. Our findings indicate that this new approach complements existing techniques used by other solvers, as evident in the experimental results (see virtual best and portfolio solvers in the plots and tables).

In the future, we aim to extend our approach to other theories such as finite fields and bit-vectors, and conduct more comprehensive experimental evaluations. Additionally, we plan to integrate this approach with different caching schemes, including value and target caches, along with periodic recaching in MCSat [23].

Acknowledgments. This material is based upon work supported in part by NSF grant 2016597. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the US Government or NSF.

The work of Enrico Lipparini was partially supported by project PRIN 2022 DeLiCE (F53D23009130001) under the MUR National Recovery and Resilience Plan funded by the European Union – NextGenerationEU.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

- Haniel Barbosa, Clark W. Barrett, Martin Brain, Gereon Kremer, Hanna Lachnitt, Makai Mann, Abdalrhman Mohamed, Mudathir Mohamed, Aina Niemetz, Andres Nötzli, Alex Ozdemir, Mathias Preiner, Andrew Reynolds, Ying Sheng, Cesare Tinelli, and Yoni Zohar. cvc5: A versatile and industrial-strength SMT solver. In Dana Fisman and Grigore Rosu, editors, Tools and Algorithms for the Construction and Analysis of Systems 28th International Conference, TACAS 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022, Proceedings, Part I, volume 13243 of Lecture Notes in Computer Science, pages 415-442. Springer, 2022.
- 2. Clark Barrett, Pascal Fontaine, and Cesare Tinelli. The Satisfiability Modulo Theories Library (SMT-LIB). www.SMT-LIB.org, 2016.
- Clark W. Barrett, Roberto Sebastiani, Sanjit A. Seshia, and Cesare Tinelli. Satisfiability Modulo Theories. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, Handbook of Satisfiability - Second Edition, volume 336 of Frontiers in Artificial Intelligence and Applications, pages 1267–1329. IOS Press, 2021.
- 4. Nikolaj Bjørner and Lev Nachmanson. Arithmetic solving in z3. In Arie Gurfinkel and Vijay Ganesh, editors, *Computer Aided Verification*, pages 26–41, Cham, 2024. Springer Nature Switzerland.
- Shaowei Cai, Bohan Li, and Xindi Zhang. Local Search for SMT on Linear Integer Arithmetic. In Sharon Shoham and Yakir Vizel, editors, *Computer Aided Verification*, pages 227–248, Cham, 2022. Springer International Publishing.
- 6. Shaowei Cai, Bohan Li, and Xindi Zhang. Local search for satisfiability modulo integer arithmetic theories. ACM Trans. Comput. Logic, 24(4), July 2023.
- Shaowei Cai and Xindi Zhang. Deep cooperation of CDCL and local search for SAT. In Chu-Min Li and Felip Manyà, editors, *Theory and Applications of Satisfiability Testing - SAT 2021*, pages 64–81, Cham, 2021. Springer International Publishing.
- 8. Shaowei Cai, Xindi Zhang, Mathias Fleury, and Armin Biere. Better decision heuristics in CDCL through local search and target phases. J. Artif. Int. Res., 74, September 2022.
- Alessandro Cimatti, Alberto Griggio, Ahmed Irfan, Marco Roveri, and Roberto Sebastiani. Experimenting on solving nonlinear integer arithmetic with incremental linearization. In Olaf Beyersdorff and Christoph M. Wintersteiger, editors, *Theory* and Applications of Satisfiability Testing – SAT 2018, pages 383–398, Cham, 2018. Springer International Publishing.
- Alessandro Cimatti, Alberto Griggio, Ahmed Irfan, Marco Roveri, and Roberto Sebastiani. Incremental linearization for satisfiability and verification modulo nonlinear arithmetic and transcendental functions. ACM Trans. Comput. Logic, 19(3), aug 2018.
- Alessandro Cimatti, Alberto Griggio, Bastiaan Schaafsma, and Roberto Sebastiani. The MathSAT5 SMT Solver. In Nir Piterman and Scott Smolka, editors, *Proceedings of TACAS*, volume 7795 of *LNCS*. Springer, 2013.

Boosting MCSat Modulo Nonlinear Integer Arithmetic via Local Search

- George E. Collins. Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In H. Brakhage, editor, Automata Theory and Formal Languages, pages 134–183, Berlin, Heidelberg, 1975. Springer Berlin Heidelberg.
- Florian Corzilius, Gereon Kremer, Sebastian Junges, Stefan Schupp, and Erika Ábrahám. SMT-RAT: An open source C++ toolbox for strategic and parallel SMT solving. In SAT, 09 2015.
- 14. Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient SMT solver. In Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS'08/ETAPS'08, page 337-340, Berlin, Heidelberg, 2008. Springer-Verlag.
- Leonardo de Moura and Dejan Jovanovic. A model-constructing satisfiability calculus. In Roberto Giacobazzi, Josh Berdine, and Isabella Mastroeni, editors, Intl. Conference on Verification, Model Checking, and Abstract Interpretation (VM-CAI), volume 7737 of LNCS, pages 1–12. Springer, 2013.
- Leonardo Mendonça de Moura and Nikolaj S. Bjørner. Satisfiability modulo theories: introduction and applications. Commun. ACM, 54(9):69-77, 2011.
- Bruno Dutertre. Yices 2.2. In Armin Biere and Roderick Bloem, editors, Computer-Aided Verification (CAV'2014), volume 8559 of Lecture Notes in Computer Science, pages 737-744. Springer, July 2014.
- Andreas Fröhlich, Armin Biere, Christoph Wintersteiger, and Youssef Hamadi. Stochastic local search for Satisfiability Modulo Theories. Proceedings of the AAAI Conference on Artificial Intelligence, 29(1), Feb. 2015.
- Zhoulai Fu and Zhendong Su. XSat: A fast floating-point satisfiability solver. In CAV, volume 9780 of Lecture Notes in Computer Science, pages 187–209. Springer, 2016.
- Carsten Fuhs, Jürgen Giesl, Aart Middeldorp, Peter Schneider-Kamp, René Thiemann, and Harald Zankl. SAT solving for termination analysis with polynomial interpretations. In João Marques-Silva and Karem A. Sakallah, editors, *Theory* and Applications of Satisfiability Testing - SAT 2007, pages 340-354, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg.
- Stéphane Graham-Lengrand, Dejan Jovanovic, and Bruno Dutertre. Solving bitvectors with MCSAT: explanations from bits and pieces. In Nicolas Peltier and Viorica Sofronie-Stokkermans, editors, Intl. Joint Conf. on Automated Reasoning (IJ-CAR), Part I, volume 12166 of LNCS, pages 103–121. Springer, 2020.
- 22. Alberto Griggio, Quoc-Sang Phan, Roberto Sebastiani, and Silvia Tomasi. Stochastic local search for SMT: Combining theory solvers with WalkSAT. In Cesare Tinelli and Viorica Sofronie-Stokkermans, editors, *Frontiers of Combining Systems*, pages 163–178, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- 23. Thomas Hader, Ahmed Irfan, and Stéphane Graham-Lengrand. Decision heuristics in MCSat (to appear). In Computer Aided Verification, 2025.
- 24. Thomas Hader, Daniela Kaufmann, Ahmed Irfan, Stéphane Graham-Lengrand, and Laura Kovács. MCSat-based finite field reasoning in the Yices2 SMT solver (short paper). In *IJCAR* (1), volume 14739 of *Lecture Notes in Computer Science*, pages 386–395. Springer, 2024.
- Thomas Hader, Daniela Kaufmann, and Laura Kovács. SMT solving over finite field arithmetic. In LPAR, volume 94 of EPiC Series in Computing, pages 238-256. EasyChair, 2023.
- 26. Leticia Hernando, Alexander Mendiburu, and Jose Lozano. Hill-climbing algorithm: Let's go for a walk before finding the optimum. pages 1–7, 07 2018.

- 20 Lipparini et al.
- Ahmed Irfan and Stéphane Graham-Lengrand. Arrays reasoning in MCSat. In SMT@CAV, volume 3725 of CEUR Workshop Proceedings, pages 24-35. CEUR-WS.org, 2024.
- Dejan Jovanović. Solving nonlinear integer arithmetic with MCSAT. In Ahmed Bouajjani and David Monniaux, editors, Verification, Model Checking, and Abstract Interpretation, pages 330-346, Cham, 2017. Springer International Publishing.
- 29. Dejan Jovanovic, Clark Barrett, and Leonardo de Moura. The design and implementation of the model constructing satisfiability calculus. In *Intl. Conf on Formal Methods in Computer-Aided Design (FMCAD)*, pages 173–180. IEEE, 2013.
- Dejan Jovanović and Leonardo de Moura. Solving non-linear arithmetic. ACM Commun. Comput. Algebra, 46(3/4):104-105, jan 2013.
- 31. Henry A. Kautz, Ashish Sabharwal, and Bart Selman. Incomplete algorithms. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, Handbook of Satisfiability - Second Edition, volume 336 of Frontiers in Artificial Intelligence and Applications, pages 213-232. IOS Press, 2021.
- 32. Gereon Kremer, Florian Corzilius, and Erika Ábrahám. A generalised branch-andbound approach and its application in SAT modulo nonlinear integer arithmetic. In Vladimir P. Gerdt, Wolfram Koepf, Werner M. Seiler, and Evgenii V. Vorozhtsov, editors, Computer Algebra in Scientific Computing - 18th International Workshop, CASC 2016, Bucharest, Romania, September 19-23, 2016, Proceedings, volume 9890 of Lecture Notes in Computer Science, pages 315-335. Springer, 2016.
- Gereon Kremer and Erika Ábrahám. Modular strategic SMT solving with SMT-RAT. Acta Universitatis Sapientiae, Informatica, 10(1):5-25, 2018.
- Bohan Li and Shaowei Cai. Local search for SMT on linear and multi-linear real arithmetic. In 2023 Formal Methods in Computer-Aided Design (FMCAD), pages 1-10, 2023.
- 35. Haokun Li, Bican Xia, and Tianqi Zhao. Local search for solving satisfiability of polynomial formulas. In Constantin Enea and Akash Lal, editors, *Computer Aided Verification*, pages 87–109, Cham, 2023. Springer Nature Switzerland.
- 36. Enrico Lipparini. Satisfiability modulo Nonlinear Arithmetic and Transcendental Functions via Numerical and Topological methods. PhD thesis, 2024.
- 37. Enrico Lipparini, Alessandro Cimatti, Alberto Griggio, and Roberto Sebastiani. Handling polynomial and transcendental functions in SMT via unconstrained optimisation and topological degree test. In Ahmed Bouajjani, Lukáš Holík, and Zhilin Wu, editors, Automated Technology for Verification and Analysis, pages 137–153, Cham, 2022. Springer International Publishing.
- 38. Enrico Lipparini and Stefan Ratschan. Satisfiability of non-linear transcendental arithmetic as a certificate search problem. In Kristin Yvonne Rozier and Swarat Chaudhuri, editors, NASA Formal Methods, pages 472–488, Cham, 2023. Springer Nature Switzerland.
- 39. Enrico Lipparini and Stefan Ratschan. Satisfiability of non-linear transcendental arithmetic as a certificate search problem. J. Autom. Reason., 69(1), January 2025.
- Yuri V. Matiyasevich. Hilbert's tenth problem. MIT Press, Cambridge, MA, USA, 1993.
- 41. Aina Niemetz, Mathias Preiner, and Armin Biere. Propagation based local search for bit-precise reasoning. *Formal Methods in System Design*, 51(3):608-636, Dec 2017.
- 42. Robert Nieuwenhuis, Albert Oliveras, and Cesare Tinelli. Solving SAT and SAT modulo theories: From an abstract Davis-Putnam-Logemann-Loveland procedure to DPLL(T). J. ACM, 53(6):937-977, November 2006.

Boosting MCSat Modulo Nonlinear Integer Arithmetic via Local Search

21

- Knot Pipatsrisawat and Adnan Darwiche. A lightweight component caching scheme for satisfiability solvers. In Theory and Applications of Satisfiability Testing-SAT 2007: 10th International Conference, Lisbon, Portugal, May 28-31, 2007. Proceedings 10, pages 294-299. Springer, 2007.
- 44. Mathias Preiner, Hans-Jörg Schurr, Clark Barrett, Pascal Fontaine, Aina Niemetz, and Cesare Tinelli. SMT-LIB release 2024 (non-incremental benchmarks), April 2024.
- Steven Prestwich and Inês Lynce. Local search for unsatisfiability. In Armin Biere and Carla P. Gomes, editors, *Theory and Applications of Satisfiability Testing -*SAT 2006, pages 283–296, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- 46. Zhonghan Wang, Bohua Zhan, Bohan Li, and Shaowei Cai. Efficient local search for nonlinear real arithmetic. In Rayna Dimitrova, Ori Lahav, and Sebastian Wolff, editors, Verification, Model Checking, and Abstract Interpretation, pages 326-349, Cham, 2024. Springer Nature Switzerland.
- 47. Tjark Weber, Sylvain Conchon, David Déharbe, Matthias Heizmann, Aina Niemetz, and Giles Reger. The SMT competition 2015-2018. J. Satisf. Boolean Model. Comput., 11(1):221-259, 2019.
- 48. Xindi Zhang, Bohan Li, and Shaowei Cai. Deep combination of CDCL(T) and local search for satisfiability modulo non-linear integer arithmetic theory. In *Proceedings* of the IEEE/ACM 46th International Conference on Software Engineering, ICSE '24, New York, NY, USA, 2024. Association for Computing Machinery.