Interpolation and Model Checking for Nonlinear Arithmetic

Dejan Jovanović(✉) and Bruno Dutertre

SRI International, Menlo Park, USA

Abstract. We present a new model-based interpolation procedure for satisfiability modulo theories (SMT). The procedure uses a new mode of interaction with the SMT solver that we call *solving modulo a model*. This either extends a given partial model into a full model for a set of assertions or returns an explanation (a model interpolant) when no solution exists. This mode of interaction fits well into the model-constructing satisfiability (MCSAT) framework of SMT. We use it to develop an interpolation procedure for any MCSAT-supported theory. In particular, this method leads to an effective interpolation procedure for nonlinear real arithmetic. We evaluate the new procedure by integrating it into a model checker and comparing it with state-of-art model-checking tools for nonlinear arithmetic.

Keywords: Satisfiability modulo theories · Craig interpolation · Nonlinear arithmetic

1 Introduction

Craig interpolation is one of the central reasoning tools in modern verification algorithms. Verification techniques such as model checking rely on Craig interpolation [11,39] as a symbolic learning oracle that drives abstraction refinement and invariant inference. Interpolation has been studied for many fragments of first-order logic that are useful in practice, such as linear arithmetic [23], uninterpreted functions [9,37], arrays [25,38], and sets [32]. In these fragments, a typical interpolation procedure constructs interpolants by traversing the clausal proof of unsatisfiability provided by an SMT solver [26,34,41] while performing interpolation locally at proof nodes. A major missing piece in the class of fragments supported by interpolating SMT solvers is nonlinear arithmetic,[1] as the

[1] By nonlinear arithmetic we mean Boolean combination of arithmetic constraints over arbitrary-degree polynomials.

This material is based upon work supported by the Defense Advanced Research Project Agency (DARPA) and Space and Naval Warfare Systems Center, Pacific (SSC Pacific) under Contract No. N66001-18-C-4011, and the National Science Foundation (NSF) grant 1816936. 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 DARPA, SSC Pacific, or the NSF.

A. Silva and K. R. M. Leino (Eds.): CAV 2021, LNCS 12760, pp. 266–288, 2021.
https://doi.org/10.1007/978-3-030-81688-9_13

complex reasoning required for nonlinear arithmetic makes fine-grained symbolic proof generation extremely difficult.

We present an approach to interpolation that is driven by models rather than proofs. Given a pair of formulas A and B such that $A \wedge B$ is unsatisfiable, an interpolant is a formula I that is implied by A and inconsistent with B. Recent model-based decision procedures, specifically the ones developed within the MCSAT [13,28] framework for SMT, are internally naturally interpolating. But, rather than interpolating two formulas, they provide a way to interpolate a set of constraints against a partial model. We capitalize on this internal ability, and extend it so that a formula A can be checked and interpolated against a partial model (*model interpolation*). This is closely related to the ability of modern SAT solvers to perform solving modulo assumptions [17], a technique that can also been used to provide interpolation capabilities in finite-state model checking [3].

We take advantage of model interpolation to build a formula-interpolation procedure through a simple idea: we can compute an interpolant of formulas A and B by iteratively interpolating (and refuting) all models of B with model interpolants from A. We develop the interpolation procedure within the MCSAT framework. This immediately allows us to generate interpolants for any theory supported by the framework. As MCSAT provides efficient complete solvers for nonlinear real arithmetic [27,29], we develop the first complete interpolation procedure for real nonlinear arithmetic.

To show that this new interpolation procedure is an effective tool that can be used on real-world problems, we integrate it into a model checker that uses interpolation for inferring k-inductive invariants. We evaluate this model checker on a set of industrial benchmarks. Our evaluation shows that the new procedure is highly effective, both in terms of speed, and the ability to support the model checker in its quest for counter-examples and invariants.

Outline. Section 2 gives background on SMT, interpolation, and nonlinear arithmetic. Section 3 presents solving modulo a model and model interpolation, and develops the general interpolation procedure. In Sect. 4, we discuss the particular needs of nonlinear arithmetic. In Sect. 5 we evaluate our implementation on nonlinear model-checking problems. We conclude in Sect. 6 and provide future research directions.

2 Background

We assume that the reader is familiar with the usual notions and terminology of first-order logic and model theory (for an introduction see, e.g., [1]).

Nonlinear Arithmetic. As usual, we denote the ring of integers with $\mathbb{Z}$ and the field of real numbers with $\mathbb{R}$. Given a vector of variables $\boldsymbol{x}$ we denote the set of polynomials with integer coefficients and variables $\boldsymbol{x}$ as $\mathbb{Z}[\boldsymbol{x}]$. A polynomial $f \in \mathbb{Z}[\boldsymbol{y}, x]$ is of the form

$$f(\boldsymbol{y}, x) = a_m \cdot x^{d_m} + a_{m-1} \cdot x^{d_{m-1}} + \cdots + a_1 \cdot x^{d_1} + a_0,$$

where $0 < d_1 < \cdots < d_m$, and the coefficients a_i are polynomials in $\mathbb{Z}[\boldsymbol{y}]$ with $a_m \neq 0$. We call x the *top variable* and the highest power d_m is the *degree* of the polynomial f. As usual, we denote with $f^{(k)}$ the k-th derivative of f in its top variable. A number $\alpha \in \mathbb{R}$ is a *root of the polynomial* $f \in \mathbb{Z}[x]$ if $f(\alpha) = 0$.

A *polynomial constraint* C is a constraint of the form $f \triangledown 0$ where f is a polynomial and $\triangledown \in \{<, \leq, =, \geq, >\}$. If the polynomial $f = f(x)$ is univariate then we also say that C is univariate. An atom is either a polynomial constraint or a Boolean variable, and formulas are defined inductively with the usual Boolean connectives ($\wedge$, $\vee$, $\neg$). The symbols $\top$ and $\bot$ denote true and false, respectively. In addition to the basic polynomial constraints, we will also be working with extended polynomial constraints. An *extended polynomial constraint* F is of the form $x \triangledown_r \mathsf{root}(f, k, x)$ where $f \in \mathbb{Z}[\boldsymbol{y}, x]$ and $\triangledown_r \in \{<_r, \leq_r, =_r, \geq_r, >_r\}$. The semantics of this predicate is the following: Given an assignment that gives real values $\boldsymbol{v}$ to the variables $\boldsymbol{y}$, then the roots of $f(\boldsymbol{a}, x)$ can be ordered over $\mathbb{R}$. If the polynomial $f(\boldsymbol{a}, x)$ has at least k real roots and α_k is the k-th smallest root[2] then the constraint is equivalent to $x \triangledown \alpha_k$. Otherwise, the constraint evaluates to $\bot$. For example, the constraint $x < \mathsf{root}(x^2 - 2, 2, x)$ represents $x < \sqrt{2}$.

Given a formula $F(\boldsymbol{x})$ we say that a type-consistent variable assignment $M = \{\boldsymbol{x} \mapsto \boldsymbol{a}\}$ satisfies F if the formula F evaluates to $\top$ in the standard semantics of Booleans and reals. We call M a model of F and denote this with $M \vDash F$. If there is such a variable assignment, we say that F is *satisfiable*, otherwise it is *unsatisfiable*. If two models M_1 and M_2 agree on the values of their common variables, we denote the model that combines M_1 and M_2 with $M_1 \cup M_2$.

Definition 1 (Craig interpolant). *Given two formulas $A(\boldsymbol{x}, \boldsymbol{y})$ and $B(\boldsymbol{y}, \boldsymbol{z})$ such that $A \wedge B$ is unsatisfiable, a* Craig interpolant *is a formula $I(\boldsymbol{y})$ such that $A \Rightarrow I$ and $I \Rightarrow \neg B$. We call the pair (A, B) an* interpolation problem.

Model Checking. A *state-transition system* is a pair $\mathfrak{S} = \langle I, T \rangle$, where $I(\boldsymbol{x})$ is a state formula describing the initial states and $T(\boldsymbol{x}, \boldsymbol{x}')$ is a state-transition formula describing the system's evolution. Given a state formula P (*the property*), we want to determine whether all reachable states of $\mathfrak{S}$ satisfy P. If this is the case, P is an *invariant* of $\mathfrak{S}$. If P is not invariant, there is a concrete trace of the system, called a *counter-example*, that reaches $\neg P$.

The direct way to prove that a property P is an invariant of $\mathfrak{S}$ is to show that it is inductive. This requires showing that P holds in the initial states: $I \Rightarrow P$, and that it is preserved by transitions: $P(\boldsymbol{x}) \wedge T(\boldsymbol{x}, \boldsymbol{x}') \Rightarrow P(\boldsymbol{x}')$. As most invariants are not inductive, a key problem in model checking is to find am *inductive strengthening* of P, that is, a property P' such that $P' \Rightarrow P$ and P' is inductive.

[2] For example, $x^2 - 2$ has two roots. The first root $-\sqrt{2}$ is the smallest of the two and the second root is $\sqrt{2}$.

Example 1 (Cauchy–Schwarz inequality). We can frame the Cauchy–Schwarz inequality as a model-checking problem in nonlinear arithmetic. The inequality is the following

$$(\sum_{i=1}^{n} x_i y_i)^2 \leq (\sum_{i=1}^{n} x_i^2)(\sum_{i=1}^{n} y_i^2). \tag{1}$$

As shown in [21], many inequalities that involve a discrete parameter (such as n above) can be converted to model-checking problems. For inequality (1), we construct the transition system $\mathfrak{S}_{cs} = \langle I, T \rangle$ where

$$I \equiv (S_1 = 0) \wedge (S_2 = 0) \wedge (S_3 = 0),$$
$$T \equiv (S_1' = S_1 + xy) \wedge (S_2' = S_2 + x^2) \wedge (S_3' = S_3 + y^2).$$

The variables S_1, S_2, S_3 correspond to the sums in (1) in order. The two variables x and y of $\mathfrak{S}_{cs}$ model the variables x_i and y_i from (1) in each iteration of $\mathfrak{S}_{cs}$. Proving the inequality amounts to showing that property $P_{cs} \equiv (S_1^2 \leq S_2 S_3)$ is an invariant of $\mathfrak{S}_{cs}$. Property P_{cs} is not inductive on its own, but property $P_{cs}' \equiv P_{cs} \wedge (S_2 \geq 0) \wedge (S_3 \geq 0)$ is an inductive strengthening of P_{cs}.

Many modern model-checking techniques, specifically those based on SMT solving, use interpolation as a tool to automatically infer inductive invariants. In this context, an interpolant can be used to over-approximate a transition in the context of a spurious counter-example. In addition to interpolation, the recent class of techniques broadly termed *property-directed reachability* (PDR) (e.g., [24,30,33]), relies on *model generalization*, which converts a concrete counter-example state into a set of counter-examples.

Definition 2 (Generalization). *Given a formula $F(\boldsymbol{x}, \boldsymbol{y})$ such that F is true in a model M, we call a formula $G(\boldsymbol{x})$ a* generalization of M *if $G(\boldsymbol{x})$ is true in M and $G(\boldsymbol{x}) \Rightarrow \exists \boldsymbol{y} \,.\, F(\boldsymbol{x}, \boldsymbol{y})$.*

A PDR model-checking procedure for nonlinear arithmetic requires both an interpolation and a generalization procedure.

3 SMT Modulo Models and Interpolation

SMT solvers typically provide an API to assert formulas and to check the satisfiability of asserted formulas. We denote with SOLVER::ASSERT(F) the solver method that adds the formula F to the set of assertions to be checked by the solver. We denote with SOLVER::CHECK() the solver method for checking satisfiability, with the following contract.

SOLVER::CHECK(): Check satisfiability of asserted formulas A and

1. if there is a model M such that $M \vDash A$, return $\langle \textbf{sat}, M \rangle$;
2. otherwise return $\langle \textbf{unsat}, \emptyset \rangle$.

In this contract, the solver does not return any form of inconsistency certificate when the assertions are unsatisfiable.[3] We generalize the standard SMT satisfiability checking to *SMT modulo models* as follows.

SOLVER::CHECK(M_0): Check satisfiability of asserted formulas A and

1. if there is a model $M \supseteq M_0$ such that $M \vDash A$, return $\langle \textbf{sat}, M, \top \rangle$;
2. otherwise return $\langle \textbf{unsat}, \emptyset, I \rangle$ where $A \Rightarrow I$ and $M_0 \vDash \neg I$.

SMT modulo models allows one to check that a formula is satisfiable modulo a partial model M_0, by seeking a solution that extends M_0. If there is no such solution, the formula I returned as the certificate of unsatisfiability is a *model interpolant*: it is implied by the assertions and inconsistent with M_0 (i.e., I evaluates to $\bot$ in the model M_0). If we restrict ourselves to Boolean formulas, SMT modulo models reduces exactly to solving modulo assumptions [17] used in the SAT community. Although this idea is not completely new, it is the first time that it is used for interpolation in SMT, as far as we know.

3.1 Interpolation

Before diving into an approach that can support the above mode of satisfiability checking, we first show how model interpolation can be used to devise a general interpolation method.

Algorithm 1: INTERPOLATE(A, B)

```
1  S_A.assert(A) ;
2  S_B.assert(B) ;
3  I ← ⊤ ;
4  while true do
5  |  ⟨r_B, M_B⟩ ← S_B.check() ;
6  |  if r_B = unsat then
7  |  |  return ⟨unsat, I⟩
8  |  ⟨r_A, M_A, I_A⟩ ← S_A.check(M_B) ;
9  |  if r_A = sat then
10 |  |  return ⟨sat, M_A ∪ M_B⟩
11 |  I ← I ∧ I_A ;
12 |  S_B.assert(I_A)
```

Algorithm 1 shows the pseudocode of a procedure that checks satisfiability and interpolates two formulas A and B. The basic idea is simple: we enumerate

[3] Some solvers support proof generation. While proofs are fundamentally important, we are interested in certificates that can always be computed and are useful in supporting further analysis. For example, proof generation for nonlinear arithmetic is still a hard open problem.

models M_k of the formula B, and refute each model M_k with a model interpolant I_k from A. If the process converges and returns **unsat**, we collect the model interpolants and construct the final interpolant $I = \bigwedge I_k$. Each interpolant I_k is implied by A because it is a model interpolant, so $A \Rightarrow I$. Each model of B is refuted by some model interpolant I_k, and so $I \Rightarrow \neg B$. On the other hand, if the process returns **sat**, the procedure has found a common model for A and B. The procedure above is model-driven and modular, in that it checks the formulas A and B independently while only communicating models (from B to A) and model interpolants (from A to B).

Lemma 1 (Correctness). *If* INTERPOLATE(A, B) *returns* $\langle \textbf{unsat}, I \rangle$ *then* $A \wedge B$ *is unsatisfiable and* I *is an interpolant for* (A, B)*. If* INTERPOLATE(A, B) *returns* $\langle \textbf{sat}, M \rangle$ *then* $A \wedge B$ *is satisfiable and* M *is a model of both* A *and* B*.*

Note that Lemma 1 does not claim termination of the procedure. Termination depends on the ability of model interpolation to produce a finite number of model interpolants that can eliminate a potentially infinite number of models.

A naive approach to check a formula $A(\boldsymbol{x}, \boldsymbol{y})$ for satisfiability modulo a model $M_0 = \{\boldsymbol{y} \mapsto \boldsymbol{v}\}$ is to use an interpolating SMT solver. First, encode the model into a formula $F_M \equiv \bigwedge(y_i = v_i)$. If the formula $A \wedge F_M$ is satisfiable in a model M, so is A and $M \supseteq M_0$. Otherwise, we compute the interpolant I of A and F_M. This naive approach satisfies the requirements of SOLVER::CHECK(M_0), but it is limited for the following reasons. First, theories such as nonlinear arithmetic have complex models and the formula F_M can be hard to express. As an example, $x \mapsto \sqrt{2}$ can only be expressed by extending the constraint language to support algebraic numbers, or by using additional assertions such as $(x^2 = 2) \wedge (x > 0)$. More important, traditional interpolation provides no guarantees in terms of convergence of a sequence of interpolation problems. For example, as already noted in [42], $\neg F_M$ would be a valid interpolant for A and F_M. But such an interpolant only eliminates a single model and could, in general, lead to non-termination of INTERPOLATE(A, B). To tackle this issue, we require that the procedure SOLVER::CHECK() produces interpolants general enough to disallow such infinite sequences of model interpolants. We do this by adopting the convergence approach and terminology of [42] to model interpolation as follows.

Definition 3 (Model Interpolation Sequence). *Given a formula* $A(\boldsymbol{x}, \boldsymbol{y})$*, a sequence of models* (M_k) *of* $\boldsymbol{y}$*, and a sequences of formulas* (I_k) *over* $\boldsymbol{y}$*, we call* (I_k) *a* model interpolation sequence *for* A *and* (M_k) *if for all* k *it holds that*

1. M_k *is consistent with* $\bigwedge_{i<k} I_i$*;*
2. M_k *is inconsistent with* A*;*
3. I_k *is a model interpolant between* A *and* M_k*.*

Definition 4 (Finite Convergence). *We say that* SOLVER::CHECK() *has the* finite convergence property *if it does not allow infinite model interpolation sequences.*

Lemma 2 (Termination). *If* SOLVER::CHECK() *has the finite convergence property, then* INTERPOLATE(A, B) *always terminates.*

3.2 SMT Modulo Models with MCSAT

We build a procedure for solving SMT modulo models by modifying the satisfiability checking procedure of MCSAT. The MCSAT method for SMT solving was introduced in [13,28] and further extended in [27]. We give a brief overview of the MCSAT terminology and mechanics, and we describe the satisfiability procedure. We emphasize modifications to the original MCSAT procedure that are needed for solving SMT modulo models.

The architecture of an MCSAT solver consists of a core solver, an assignment trail, and reasoning plugins. The *core solver* drives the overall solving process, and is responsible for dispatching notifications and handling requests from the plugins. The *solver trail* is a chronological record that tracks assignments of terms to values. It is shared by the core solver and the reasoning plugins. The *reasoning plugins* are modules dedicated to handling specific theory terms and constraints (e.g., clauses for Booleans, polynomial constraints for arithmetic). A plugin reasons about the content of the solver trail with respect to the set of currently relevant terms. In the context of nonlinear arithmetic problems, the reasoning plugins are the arithmetic plugin and the Boolean plugin. The most important role of the core solver is to perform conflict analysis when one of the reasoning plugins detects a conflicting state.

When formulas $F_1, \ldots, F_n$ are asserted, by calling SOLVER::ASSERT(F_i), the core solver notifies all plugins of the asserted formulas. The plugins analyze the formulas and report all *relevant terms* back to the core. The relevant terms are the variables and subterms of the formulas F_is that need to be consistently assigned to ensure a satisfying assignment. In nonlinear arithmetic, relevant terms are all variables, arithmetic constraints, and non-negated Boolean terms that appear in the input formula (or are part of a learnt clause). Once the relevant terms are collected, the core solver adds the assertions to the trail. The initial trail contains then the partial assignment $F_i \rightsquigarrow \top$ and the search for a full satisfying assignment starts from this trail.

Solver Trail and Evaluation. The assignment trail is the central data structure in the MCSAT framework. It is a generalization of the Boolean assignment trail used in modern CDCL SAT solvers. The trail records a partial (and potentially inconsistent) model that assigns values to relevant terms. If the satisfiability algorithm terminates with a **sat** answer, the full satisfying assignment can be read off the trail. At any point during the search, the trail can be used to evaluate any relevant compound term based on the values of its sub-terms. A term t (and $\neg t$, if Boolean) *can be evaluated* in the trail M if t itself is assigned in M, or if all closest relevant sub-terms of t are assigned in M (and its value can therefore be computed). As the search progresses, it is possible for some terms to *be evaluated in two different ways*, which can result in a conflict (i.e., a term assigned different values). In order to account for this ambiguity, we define an evaluation predicate $\mathsf{evaluates}[M](t, v)$ that returns **true** if the term t can evaluate to the value v in trail M.

Algorithm 2: MCSAT::CHECK($\boldsymbol{x} \mapsto \boldsymbol{v}$)

```
Data: solver trail M, relevant variables/terms to assign in queue
while true do
    unitPropagate() ;
    if a plugin detected a conflict and the conflict clause is C then
        ⟨C, final⟩ ← analyzeConflict(M, C, x) ;
        if final then
            I ← analyzeFinal(M, C) ;
            return ⟨unsat, I⟩
        else backtrackWith(M, C) ;
    else
        if exists x_i ∈ x unassigned in M then
            ownerOf(x_i).decideValue(x_i, v_i)
        else
            if queue.empty() then return ⟨sat, M⟩ ;
            x ← queue.pop() ;
            if x is unassigned then ownerOf(x).decideValue(x) ;
```

Conflicts and Conflict Clauses. One of the main responsibilities of reasoning plugins is to ensure that the trail is consistent at any point in the search. A trail is *evaluation consistent* if no relevant term can evaluate to two different values, as described above. A trail is *unit consistent* if every relevant term can be given a value without making the trail evaluation inconsistent. If the trail is not evaluation consistent or unit consistent, the trail is *in conflict*.

Trail consistency is a generalization of the consistency that CDCL SAT solvers enforce during their search. By unit propagation, a SAT solver ensures that, if no conflict has been detected, no clause can be falsified by assigning a single variable (i.e., no clause evaluates to both $\top$ and $\bot$). In the MCSAT framework, the plugins do the same: they keep track of unit constraints and reason about the consistency of the trail. It is the responsibility of the plugin to report conflicts. Each conflict must be accompanied with a *valid* conflict clause that explains the inconsistency.[4] A clause $C \equiv (L_1 \vee \ldots \vee L_n)$ is a *conflict clause* in a trail M, if each literal L_i can evaluate to $\bot$ in M, i.e. if $\mathsf{evaluates}[M](L_i, \bot)$.

Example 2. Consider the constraint $C \equiv (x^2 + y^2 < 1)$ with the set of relevant terms $\{C, x, y\}$, and the following solver trails

$$M_1 = [\![\, C \mapsto \top, x \mapsto 0 \,]\!], \qquad M_2 = [\![\, C \mapsto \top, x \mapsto 0, y \mapsto 0 \,]\!],$$
$$M_3 = [\![\, C \mapsto \top, x \mapsto 1 \,]\!], \qquad M_4 = [\![\, C \mapsto \top, x \mapsto 1, y \mapsto 0 \,]\!].$$

The trails M_1 and M_2 are consistent, the trail M_3 is unit inconsistent (no consistent assignment for y exists), and M_4 is evaluation inconsistent (C evaluates to both $\top$ and $\bot$). A valid explanations for the inconsistency of M_3 is the conflict

[4] By valid here we mean that the clause is a universally true statement on its own.

clause $C_3 \equiv \neg C \vee (x < 1)$, while a valid explanation for the inconsistency of M_4 is the conflict clause $C_4 \equiv \neg C \vee C$. Although C_4 is a tautology, it is an acceptable conflict clause since both literals can evaluate to $\bot$ (because $\mathsf{evaluates}[M_4](C, \top)$ and $\mathsf{evaluates}[M_4](C, \bot)$).

Main Procedure. The implementation of the satisfiability checking procedure SOLVER::CHECK() is a generalization of the search-and-resolve loop of modern SAT solvers (see, e.g. [16,17]). The procedure is shown in Algorithm 2, where we emphasize the extensions needed for SMT modulo models in red. The overall procedure performs a direct search for a satisfying assignment and terminates either by finding an assignment that extends the given partial model, or deduces that the problem is unsatisfiable as certified by an appropriate model interpolant.

The main elements of the procedure are unit propagation and decisions, used for constructing the assignment, and conflict analysis for repairing the trail when it becomes inconsistent. The `unitPropagate()` procedure invokes the propagation procedures provided by the plugins. Propagation allows each plugin to add new assignments to the top of the trail. If, during propagation, a plugin detects an inconsistency, it reports the conflict to the core solver along with a valid conflict clause. The `decideValue(`x`)` procedure assigns a value of the given unassigned term x. Decisions are performed only after propagation has fully saturated with no reported conflicts, which means that the trail is unit consistent. In such a trail, an assignment for x is guaranteed to exist, but the choice of a particular value is delegated to the plugin responsible for x (e.g., the arithmetic plugin for real-typed terms).

Modification 1 (Decisions). *To support SMT modulo a model $\boldsymbol{x} \mapsto \boldsymbol{v}$, variables $x_i \in \boldsymbol{x}$ of the input model are decided before any other term, and are assigned the provided value v_i. The procedure that performs this decision is denoted with* `decideValue(`x_i`,` v_i`)`*. If a decision introduces an evaluation inconsistency, the plugin reports the conflict with a conflict clause.*

Detecting and explaining decision conflicts is straightforward: there must exist a single constraint C that can evaluate to both $\top$ and $\bot$ in the trail. Such conflicts can always be explained with a clause of the form $(\neg C \vee C)$.

If a conflict is reported, either during propagation or in a decision, the procedure invokes the conflict analysis procedure `analyzeConflict()`. This procedure takes the reported conflict clause C and finds the root cause of the conflict. The analysis backtracks the trail, element by element, so long as C is a conflict clause, while resolving any trail propagations from C. Once done, the analysis returns the clause along with the flag that indicates whether this conflict clause C is empty (indicating the final conflict). If the conflict is not final, the procedure calls `backtrackWith()` to backtrack the trail further, if possible, and add a new assignment to the trail, ensuring progress and fixing the conflict. The main invariant of the conflict resolution procedure is that the *conflict clause C is always implied by asserted formulas.*

Modification 2 (Conflict Analysis). *To support SMT modulo a model* $\boldsymbol{x} \mapsto \boldsymbol{v}$*, the analysis procedure* `analyzeConflict(`M, C, $\boldsymbol{x}$`)` *stops as soon as it encounters a variable* $x_i \in \boldsymbol{x}$ *to resolve, and returns* $\langle C, \textbf{true} \rangle$.

This modification is based on the fact that the variables x_i have a fixed value given by the model. Assume that conflict analysis attempts to undo a variable x_i that is part of the provided model $\boldsymbol{x} \mapsto \boldsymbol{v}$. This can only happen when the trail consists of only variables from $\boldsymbol{x}$ and implications of asserted formulas. In other words, this particular conflict cannot be resolved unless we modify either the assertions themselves or the input model. The clause resulting from the analysis marked as final is our starting point for producing the model interpolant.

Modification 3 (Final Analysis). *To support SMT modulo a model* $\boldsymbol{x} \mapsto \boldsymbol{v}$*, the procedure* `analyzeFinal(`$M$, C`)` *resolves any remaining trail propagations in* M *from the clause* C *and returns the resulting clause* I.

The resolution of propagations in this final analysis is done in the same manner as in regular conflict analysis. This means that the resulting clause I is implied by the asserted formulas. In addition, resolving all propagations from the conflict clause ensures that all literals of I evaluate to false only because of the assignment $\boldsymbol{x} \mapsto \boldsymbol{v}$, making I an appropriate model interpolant.

Example 3. Consider two formulas $F_1 \equiv b$ and $F_2 \equiv \neg b \vee (x^2 + y^2 < 2)$. When asserting these two formulas to the MCSAT solver, the Boolean and arithmetic plugins will identify the set of terms relevant for satisfiability as $R = \{b, x, y, (x^2 + y^2 < 2)\}$. Additionally, the assertions will be added to the trail and propagated[5], resulting in the following initial trail

$$M_0 = [\![\, b \rightsquigarrow \top, F_2 \rightsquigarrow \top, (x^2 + y^2 < 2) \overset{F_2}{\rightsquigarrow} \top \,]\!].$$

We now apply our procedure to solve F_1 and F_2 modulo the partial model $\{x \mapsto 2\}$.

In the first iteration, no term in R is unit (with only one variable unassigned), and propagation does not infer any new facts or conflicts. The procedure thus perform a decision on the unassigned variable x of the model, resulting in the trail

$$M_1 = [\![\, b \rightsquigarrow \top, F_2 \rightsquigarrow \top, (x^2 + y^2 < 2) \overset{F_2}{\rightsquigarrow} \top, x \mapsto 2 \,]\!].$$

In the second iteration, as $(x^2 + y^2 < 2)$ is unit in the trail M_1, the arithmetic plugin examines the constraint and deduces that there is no potential solution for y. This constitutes a unit inconsistency that the plugin reports, along with the conflict clause[6]

$$C_0 \equiv \neg(x^2 + y^2 < 2) \vee \neg(x > \sqrt{2}).$$

[5] Notation $t \overset{F}{\rightsquigarrow} v$ denotes that t is assigned to v due to propagation, and F is the reason of the propagation.

[6] We use $(x > \sqrt{2})$ as a shorthand for the extended constraint $x >_r \mathsf{root}(x^2 - 2, 2, x)$.

Conflict analysis takes clause C_0 and starts the resolution process. As the top variable x on the trail M_1 is part of the input model, the analysis stops and reports that the clause C_0 is the final explanation. This clause is valid, but not yet a model interpolant as it contains a literal with variable y. We then proceed with the final analysis to remove such literals. First, we resolve $(x^2 + y^2 < 2)$ from C_0 using its reason clause F_1, which gives the clause $C_1 \equiv \neg b \vee \neg(x > \sqrt{2})$. Then, we resolve b from C_1 with an empty reason (b is an assertion), resulting in the final clause and model interpolant $I = \neg(x > \sqrt{2})$.

4 Nonlinear Arithmetic

The general approach to interpolation presented so far is not specific to nonlinear arithmetic. We now tackle two practical issues that arise in nonlinear arithmetic and we discuss the properties of our interpolation procedure in the context of nonlinear arithmetic. First, on nonlinear problems, as seen in Example 3, the interpolation procedure can return model interpolants that include extended polynomial constraints. This is an artifact of the underlying decision procedure (such as NLSAT [29]) that might use extended polynomial constraints to succinctly represent conflict explanations. While such constraints make decision procedures more effective, they are undesirable for interpolation: interpolants should be described in the language of the input formulas, if possible. Second, to use the interpolant procedure in the context of model checking, we also need to devise a generalization procedure for polynomial constraints.

This section uses concepts from cylindrical algebraic decomposition (CAD). We keep the presentation example-driven and focused on our particular needs, and refer the reader to the existing literature for further information [2,5,7]. Cylindrical algebraic decomposition is a general approach for reasoning about polynomials based on the following result due to Collins [10]. For any set of polynomials $f_1, \ldots, f_k \in \mathbb{Z}[x_1, \ldots, x_n]$ one can algorithmically decompose $\mathbb{R}^n$ into connected regions (called cells) such that all the polynomials f_j are sign-invariant in every cell C_i. This means that the cells also maintain the truth value of any polynomial constraints over the polynomials f_i, which is crucial in many reasoning techniques for polynomial constraints.

The theory and practice of CAD is heavily dependent on the ordering of variables involved. For this paper we always assume the CAD order to be the same as the order of the defined polynomials (e.g., $x_1 < x_2 < \ldots < x_n$). Every CAD cell is cylindrical in nature, and can be described by constraints where every dimension of the cell (called a level) can be completely defined by relying only on the previous dimensions. We illustrate this through an example.

Example 4. Consider the polynomial $f = x^2 + y^2 - 2 \in \mathbb{Z}[x, y]$. A CAD of f is depicted in Fig. 1 (left). The cell C_1 is defined by two constraints:

$$C_1^y \equiv y >_r \mathsf{root}(x^2 + y^2 - 2, 2, y),$$
$$C_1^x \equiv x >_r \mathsf{root}(x^2 - 2, 1, x) \wedge x <_r \mathsf{root}(x^2 - 2, 2, x).$$

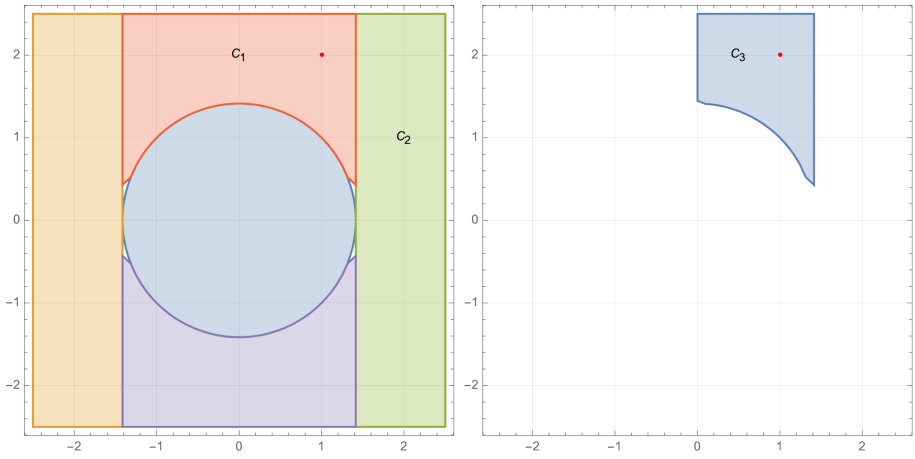

Fig. 1. CAD of the polynomial $f = x^2 + y^2 - 2$ from Example 4 (left). Computed cell capturing the model $(1, 2)$ of Example 5 (right).

Constraint C_1^x is at the first level (it's a constraint on x only), while constraint C_1^y is at the second level and relates variables x and y. The full cell description is then $C_1 \equiv C_1^x \wedge C_1^y$. The green cell C_2 can be described by $C_2^y \equiv \top$ and $C_2^x \equiv x >_r \mathsf{root}(x^2 - 2, 2, x)$, with the full description $C_2 \equiv C_2^y \wedge C_2^x$.

Model-based decision procedures such as NLSAT rely on CAD construction but do not construct the complete CAD decomposition. Instead, given a point in $\mathbb{R}^n$ they can construct a single cell of a CAD in a model-driven fashion. For more information about this approach, we refer the reader to [4,27]. For our purposes we abstract the cell construction, and denote with $\mathsf{describeCell}(F, M)$ the function that, given a set of polynomials F, returns a description of a CAD cell of F that contains the model M.

Following the terminology used in CAD, we say that a non-empty connected subset of $\mathbb{R}^k$ is a *region*. A set of polynomials $\{f_1, \ldots f_s\} \subset \mathbb{Z}[\boldsymbol{y}, x]$, with $\boldsymbol{y} = \langle y_1, \ldots, y_n \rangle$, is said to be *delineable* in a region $S \subseteq \mathbb{R}^n$ if for every f_i (and f_j) from the set, the following properties are invariant for any $\boldsymbol{\alpha} \in S$:

1. the *total number of complex roots* of $f_i(\boldsymbol{\alpha}, x)$;
2. the *number of distinct complex roots* of $f_i(\boldsymbol{\alpha}, x)$;
3. the *number of common complex roots* of $f_i(\boldsymbol{\alpha}, x)$ and $f_j(\boldsymbol{\alpha}, x)$.

Delineability has important consequences on the number and arrangement of real roots of polynomials f_i. As explained by the following theorem, if a set of polynomials F is delineable on a region S, then the number of real roots of the polynomials does not change on S. Moreover, these roots maintain their relative order on the whole of S.

Theorem 1 (Corollary 8.6.5 of [40]). *Let F be a set of polynomials in $\mathbb{Z}[\boldsymbol{y}, x]$, delineable in a region $S \subset \mathbb{R}^n$. Then, the real roots of F vary continuously over S, while maintaining their order.*

For a polynomial $f \in \mathbb{Z}[\boldsymbol{x}]$ and model $M = \{\boldsymbol{x} \mapsto \boldsymbol{v}\}$, we denote with $\mathsf{sgncstr}(f, M)$ the polynomial constraint that matches the sign of f in M, i.e.

$$\mathsf{sgncstr}(f, M) = \begin{cases} f < 0 & \text{if } \mathsf{sgn}(f(\boldsymbol{v})) < 0 \\ f > 0 & \text{if } \mathsf{sgn}(f(\boldsymbol{v})) > 0 \\ f = 0 & \text{if } \mathsf{sgn}(f(\boldsymbol{v})) = 0 \end{cases}$$

As described above, a CAD cell can be succinctly described by relying on extended polynomial constraints. We now show that the description of the cell can be reduced to basic polynomial constraints.

Lemma 3. *Let $f_i \in \mathbb{Z}[y_1, \ldots, y_n, x]$ be two polynomials of degrees m_i, and $F_i \equiv x \bigtriangledown_r \mathsf{root}(f_i, k_i, x)$ be extended polynomial constraints of a cell description. Let S be a region of $\mathbb{R}^n$ where $\{f_1, f_2\}$ are delineable and let $M = \{\boldsymbol{y} \mapsto \boldsymbol{v}, x \mapsto \alpha\}$ be a model such that $\boldsymbol{v} \in S$. Then, for all $\boldsymbol{y} \in S$ it holds that*

$$\bigwedge_{i=0}^{m_1-1} \mathsf{sgncstr}(f_1^{(i)}, M) \wedge \bigwedge_{i=0}^{m_2-1} \mathsf{sgncstr}(f_2^{(i)}, M) \Rightarrow F_1 \wedge F_2.$$

The proof of this lemma is relatively straightforward. The CAD cell description for level x represents an entry in the sign table of f_1 and f_2 (with no roots in between). A part of this sign table entry that contains M can be described with the signs of all the derivatives of f_1 and f_2 as long as we can guarantee that neither the arrangement nor the number of roots f_1 and f_2 change. But, this is guaranteed by f_1 and f_2 being delineable on S, so the lemma holds.

As a corollary to this lemma, in the context of CAD cell construction around a model M, we can replace any extended constraints describing a cell C with basic constraints stating that the signs of the polynomial derivatives are the same as in M. This results in a valid CAD subcell $C' \subseteq C$ for the same polynomials, that still contains the model M. We denote the function that constructs a basic CAD cell description of a set of polynomials F capturing the model M with $\mathsf{describeCellBasic}(F, M)$.

Example 5. Based on Example 4, we can construct a cell around the model $M = \{x \mapsto 1, y \mapsto 2\}$. Function $\mathsf{describeCellBasic}(F, M)$ will return the constraints

$$C_3^y \equiv (x^2 + y^2 > 2) \wedge (y > 0),$$
$$C_3^x \equiv (x^2 < 2) \wedge (x > 0).$$

The full cell description is then $C_3 \equiv C_3^x \wedge C_3^y$. Note that this cell is smaller than the cell C_1 from Example 4. This reduction in size is generally undesirable, but it is a price to pay for having the description in a simpler language.

Interpolation Without Extended Constraints. We now show how the cell construction described above can be used to remove extended polynomial constraints from a model interpolant. Assume a clausal model interpolant

$$I = (L_1 \vee \ldots \vee L_i \vee \ldots \vee L_N)$$

that is implied by formula A and refutes a model $M = \{\boldsymbol{x} \mapsto \boldsymbol{v}\}$, i.e., all literals of I evaluate to $\bot$ in M. Assume also that some literal L_i contains an extended polynomial constraint $x_n \bigtriangledown_r \mathsf{root}(f, k, x_n)$, with $f \in \mathbb{Z}[\boldsymbol{x}]$. We aim to replace the extended literal L_i with literals over basic polynomial constraints. To do so, we need to find literals $L_i^1, \ldots, L_i^m$ such that $L_i \Rightarrow (L_i^1 \vee \ldots \vee L_i^m)$ and all literals L_i^j evaluate to $\bot$ in M. Then, the clause

$$I' = (L_1 \vee \ldots \vee L_i^1 \vee \ldots L_i^m \vee \ldots \vee L_N)$$

will also be a model interpolant implied by A that refutes the model M.

We can construct the literals L_i^j using single cell construction as follows. We create a description of the CAD cell of the polynomial f from L_i that captures the model M. Let $\mathsf{describeCellBasic}(\{f\}, M) = D_1 \wedge \ldots \wedge D_m$ be this description. Since the cell fully captures the behavior of f around M, we know that $D_1 \wedge \ldots \wedge D_m \Rightarrow \neg L_i$ and all literals D_j evaluate to $\top$. Therefore, we can use the cell description to eliminate the extended literal L_i, obtaining the clause

$$I' = (L_1 \vee \ldots \vee \neg D_i \vee \ldots \neg D_m \vee \ldots \vee L_n)$$

By continuing this process, we can replace all extended literals from a model interpolant, to obtain a model interpolant in the basic language of polynomial constraints.

Example 6. Consider the model interpolant $I = \neg(x >_r \mathsf{root}(x^2 - 2, 2, x)$ from Example 3 that refutes the model $M = \{x \mapsto 2\}$. To express I in terms of basic polynomials constraints we first construct a regular CAD cell of $f = x^2 - 2$ around M. In this case this cell is simply $x >_r \mathsf{root}(x^2 - 2, 2, x)$. Then, we use Lemma 3 to construct a basic CAD cell description as $(x^2 > 2) \wedge (x > 0)$. Finally, the simplified interpolant is $I' = \neg(x^2 > 2) \vee \neg(x > 0)$.

Termination. With the description of the interpolation procedure complete, we discuss the termination of the procedure. To do so, we fix the formula $A(\boldsymbol{x}, \boldsymbol{y})$ of Definition 3 and we assume a fixed order of variables that ensures $y_i < x_i$. Since the MCSAT decision procedure on which we rely is based on CAD, we can put a bound on the set of literals that can ever appear in a model interpolant from the formula A to an arbitrary model M. Let P_A be the set of polynomials appearing in A, and let $P = \mathsf{P}(P_A)$ denote the closure of the set P_A under the CAD projection operator used by the decision procedure. Finally, let P' be the closure of P under derivatives. The set of polynomial constraints that can appear in the interpolant I is limited to basic polynomial constraints over polynomials in P'. This means that the procedure MCSAT::CHECK() can only generate a finite number of model interpolants and therefore has the finite convergence property.

Lemma 4. *Assuming a fixed variable order, the* MCSAT::CHECK() *procedure has the finite convergence property for nonlinear arithmetic formulas.*

Together with Lemma 2, this lemma implies that our interpolation procedure for the theory of nonlinear arithmetic terminates.

Model Generalization. We now proceed to show how the CAD cell construction can be used in a natural way to provide model-driven generalization. As in Definition 2, assume a formula $F(\boldsymbol{x}, \boldsymbol{y})$ such that F is true in a model M. Our aim is to construct a formula $G(\boldsymbol{x})$ that generalizes the model M and still guarantees a solution to F.

Following the approach of [15], we do this in two steps. First, we construct an implicant B of F based on the model M. Then, we eliminate the variables $\boldsymbol{y}$ from B, again relying on the model M. The implicant B is a conjunction of literals that implies F and such that B is true in M. The implicant can be computed by a top-down traversal of the formula F while using the model M to evaluate the formula nodes (see, e.g., [15] for a detailed description). To find a formula G such that $G \Rightarrow \exists \boldsymbol{y} \,.\, B$, we use CAD cell construction as follows. Let $P \subseteq \mathbb{Z}[\boldsymbol{x}, \boldsymbol{y}]$ be the set of all polynomials appearing in B, and let the cell description of P around M be

$$\mathsf{describeCellBasic}(P, M) = D_{\boldsymbol{x}} \wedge D_{\boldsymbol{y}}.$$

Here, $D_{\boldsymbol{x}}$ denotes the description of cell levels of variables $\boldsymbol{x}$, while $D_{\boldsymbol{y}}$ denotes the description of cell levels of variables $\boldsymbol{y}$. Because of the cylindrical nature of CAD cells, and the order on variables y_i and x_i, we are guaranteed that every solution of $D_{\boldsymbol{x}}$ can be extended to a solution of $D_{\boldsymbol{y}}$. Therefore we set the final generalization $G(x) \equiv D_{\boldsymbol{x}}$.

Example 7 (Generalization). Consider the formula $F \equiv (x^2 + y^2 < 2)$ and the model $M = \{x \mapsto 1, y \mapsto 2\}$ that satisfies F, and let us compute a generalization $G(x)$ of M. First, we compute a CAD cell of $f = x^2 + y^2 - 2$ as shown in Example 5. Then we drop the description of cell level y, to obtain the model generalization $G(x) \equiv (x^2 < 2) \wedge (x > 0)$.

5 Evaluation

To the best of our knowledge, there is no clear metric for evaluating how good an interpolant is, or for comparing different interpolants. In this section, we first show two examples to illustrate the procedure and its applications. Then, we evaluate the effectiveness of our interpolation procedure on practical problems that arise from model-checking applications. To this end, we integrate the procedure into a model checker and evaluate whether the procedure is efficient, and can produce abstractions that help the model checker synthesize invariants and discover counter-examples.

We have implemented the reasoning procedures (solving modulo partial models and interpolation procedure) by extending the existing MCSAT implementation of the YICES2 SMT solver [14]. We used the LIBPOLY library [31] for computing the model generalization and simplification of algebraic cells. Since YICES2 is integrated into the SALLY model checker [30], we rely on the PDKIND method [30] as the model checking engine (the user of interpolation) in our evaluation.

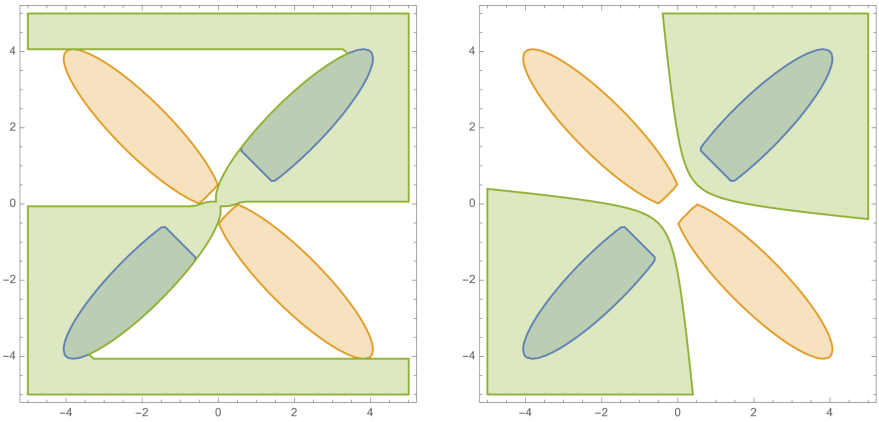

Fig. 2. Illustration of interpolants from Example 8. In blue and orange are the feasible space of the formulas A and B (projected on x and y). In green is the feasible space of the interpolant produced by our method (on the left) and the interpolant produced by [19] (on the right). (Color figure online)

Example 8. We compare the style of interpolants generated by our new procedure with the ones generated by numerical approaches such as [19]. Example 4 from [19] considers two formulas of the form

$$A(x, y, a_1, a_2, b_1, b_2) \equiv (f_1 \geq 0 \wedge f_2 \geq 0) \vee (f_3 \geq 0 \wedge f_4 \geq 0),$$
$$B(x, y, c_1, c_2, d_1, d_2) \equiv (g_1 \geq 0 \wedge g_2 \geq 0) \vee (g_3 \geq 0 \wedge g_4 \geq 0).$$

The polynomials f_i and g_i involved in A and B are of degree 2. The right-hand side of Fig. 2 shows the interpolant I_1 found by the approach in [19]. This interpolant is of the form $h(x, y) > 0$, where h is a polynomial degree two computed using semidefinite programming. Our approach, on the other hand, produces the interpolant I_2 shown on the left-hand side of Fig. 2. This interpolant consists of 12 clauses, each containing 6–8 polynomial constraints over 16 different polynomials (8 linear, 8 of degree 2). The interpolant I_2 is ultimately produced from fragments of a CAD so its edges touch upon the critical points of the shape they were produce from (formula A). Interpolant I_1, on the other hand, has a simple form dictated by the method [19]. Which form is ultimately more useful depends on a particular application.

	IC3-NRA			KIND			PDKIND		
problem set	solved	valid/invalid	time (s)	solved	valid/invalid	time (s)	solved	valid/invalid	time (s)
handcrafted (14)	10	9/1	381	3	2/1	0	**14**	13/1	4
hycomp (7)	2	2/0	15	4	1/3	796	**4**	2/2	792
hyst (65)	**39**	32/7	404	25	13/12	50	38	26/12	42
isat3 (1)	0	0/0	0	0	0/0	0	0	0/0	0
isat3-cfg (10)	8	6/2	14	9	6/3	9	**10**	7/3	8
nuxmv (2)	**2**	2/0	158	0	0/0	0	1	1/0	1118
sas13 (13)	10	5/5	13	5	0/5	0	**13**	8/5	7
tcm (2)	2	2/0	1	**2**	2/0	0	**2**	2/0	0
	73	58/15	986	48	24/24	855	**82**	59/23	1971

Fig. 3. Evaluation Results. For each tool, we report the number of solved problems, how many of the solved problems were valid and invalid, and the total time used to solve them. The rows correspond to different problem classes, and the bottom row reports the overall results for all 114 benchmarks.

Example 9 (Cauchy-Schwartz). As described in Example 1, we can model the computation of Cauchy-Schwarz inequality as a transition system $\mathfrak{S}_{cs}$. Then we can prove the inequality correct if we can prove that the property P_{cs} is valid in $\mathfrak{S}_{cs}$. The PDKIND model checking engine with the new interpolation procedure proves the property valid in 1 s.

Benchmarks. We run the evaluation on an existing set of nonlinear model-checking problems used by Cimatti, et al. [8]. This set consists of 114 benchmarks from various sources: handcrafted benchmarks, hybrid system verification, NUXMV benchmarks, C floating-point verification, and verification of Simulink models. The benchmark problems all contain transition systems with nonlinear behavior. For each problem, the goal is to prove or disprove a single invariant. We refer the reader to [8] for a more detailed description.

Evaluation. Cimatti, et al. [8] present an abstraction approach based on incrementally more precise linear approximations of nonlinear polynomials. They show that this approach, implemented in the IC3-NRA tool, is superior to other tools (such as, ISAT3 [36] and NUXMV [6] with upfront linear abstraction). Since our goal is to show the effectiveness of our interpolation procedure, rather than compare to many model checking engines, we keep the evaluation simple and only compare to IC3-NRA. In addition, we include the k-induction engine KIND of SALLY in the comparison to illustrate the importance of invariant inference and counter-example generation.[7]

We ran the tools on the benchmark set with a 1 h CPU timeout per problem. The results are shown in Fig. 3 and on the cactus plot in Fig. 4. A scatter plot comparison of PDKIND against IC3-NRA and KIND is shown in Fig. 5.

[7] KIND performs k-induction checks for increasing values of k and stops if either the property is shown k-inductive, or a counter-example is found.

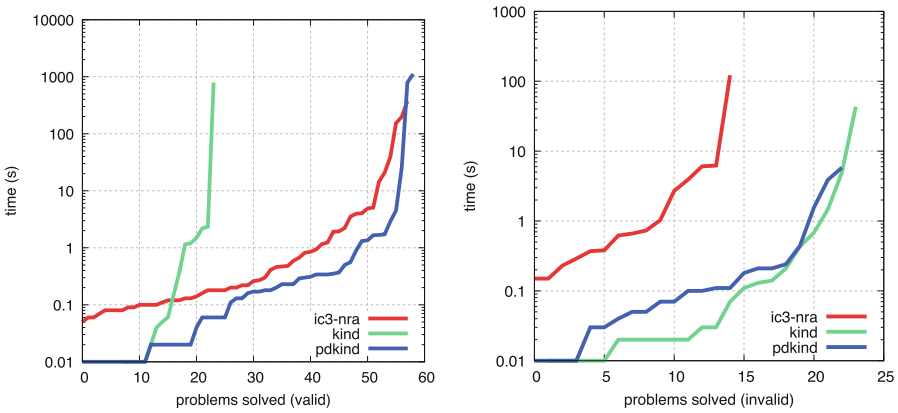

Fig. 4. Cactus Plots Comparing the Performance of IC3-NRA, KIND, and PDKIND. The x axis is the number of problems solved (valid on the left, invalid on the right) and the y axis is the time needed to solve the problem (log scale).

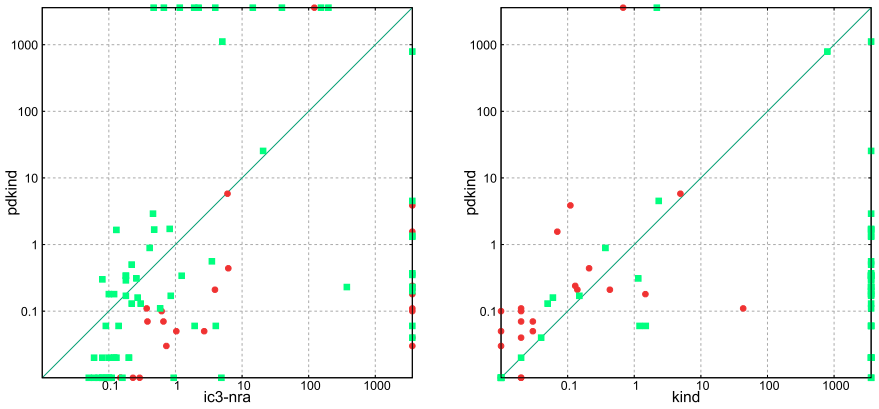

Fig. 5. Scatter Plots Comparing the Performance of IC3-NRA and KIND with PDKIND. Green squares represent problems that are valid. Red dots represent problems that are invalid. Each axis represents the time it took the tool to solve the problem (log scale). (Color figure online)

As can be seen from Fig. 3, the results are positive. The PDKIND engine with the new interpolation method can prove more properties and find more counter-examples than the state-of-the-art IC3-NRA.

Out of 59 properties that PDKIND shows correct, 36 cannot be proved by KIND. This means that these properties are likely not k-inductive and that the interpolants produced by our procedure are valuable abstractions in invariant inference. Similarly, IC3-NRA proves 37 properties that are not k-inductive. As can be seen from the scatter plot in Fig. 5, there are properties that PDKIND can prove than IC3-NRA cannot, and vice versa (11 and 10, respectively). This is to

be expected from a difficult domain, but it also means that the interpolation and the abstraction approach (or other methods) can be used to complement each other.

As for the invalid properties, since our interpolation method (and thus PDKIND) is based on complete and precise reasoning, while IC3-NRA relies on abstraction, it is to be expected that PDKIND can prove more properties invalid. Furthermore, the comparison with KIND in Fig. 5 shows that PDKIND finds all but one counter-examples that KIND does in a similar amount of time. We see this as a confirmation that the interpolation and generalization methods are effective, i.e., they do not impede the search for counter-examples.

5.1 Related Work

There is ample literature on interpolation for different fragments of nonlinear arithmetic. Existing methods can roughly be classified into two categories: approaches based on interval reasoning, and approaches based on semidefinite programming. Interval reasoning techniques (e.g., [20,35,36]) construct a proof of unsatisfiability through interval slicing and propagation. From such a proof, interpolants can be built using proof-based interpolation techniques. While incomplete, interval-based techniques can be very effective on problems that are hard for complete techniques. Moreover they can support more polynomial functions (e.g., elementary functions, ODEs). Our procedure is complete, but it is limited to the theories supported by MCSAT. The approaches based on semidefinite programming [12,18,19] generally approach the interpolation problem by restricting both the fragment of arithmetic (e.g., bounded constraints, same set of variables, quadratic constraints) and the shape of the interpolant (a single polynomial constraint) so that the interpolant itself can be represented as a semidefinite optimization problem. When they apply, these procedures are also very effective but they suffer from numerical imprecision, requiring special care to account for these errors and making them difficult to use in formal verification. In contrast, out procedure applies to nonlinear arithmetic as a whole. It relies on symbolic techniques, which are not subject to numerical errors. It is precise and complete, and it produces clausal interpolants.

The core ideas beyond our model-based interpolation approach were presented at the Boolean level as SAT solving with assumptions [17]. Closest to our work is the work of Schindler and Jovanović [42] where a similar model-based approach to interpolation is applied to conjunctions of linear arithmetic constraints based on conflict resolution. Our work is more general as it applies to formulas other than conjunctions, and it is applicable to a wider range of theories.

6 Conclusion and Future Work

We have presented a general approach for interpolation in SMT. This novel approach relies on a mode of interaction with the SMT solver that can check a

formula for satisfiability modulo a partial model and, if the formula is unsatisfiable, can return a model interpolant that refutes the model. This allows us to develop a first complete interpolation procedure for nonlinear arithmetic. We have implemented the new procedure in the YICES2 SMT solver and evaluated the interpolation procedure on model-checking problems. The new procedure seems to be effective in practice and opens new possibilities in the verification of systems that contain nonlinear behavior. Additionally, we show interesting examples of how the procedure can be used in automating induction proofs in mathematics.

The interpolation procedure that we presented can support other theories available in MCSAT (e.g., uninterpreted functions [28], bit-vectors [22], nonlinear integer arithmetic [27]). We plan to explore interpolation in these theories in more detail, and in the contexts where interpolation can be beneficial (e.g., model checking, quantified reasoning, termination, and proof generation).

References

1. Barrett, C., Tinelli, C.: Satisfiability modulo theories. In: Handbook of Model Checking, pp. 305–343. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-10575-8_11
2. Basu, S., Pollack, R., Roy, M.-F.: Algorithms in Real Algebraic Geometry. Springer, Heidelberg (2006). https://doi.org/10.1007/3-540-33099-2
3. Bayless, S., Val, C.G., Ball, T., Hoos, H.H., Hu, A.J.: Efficient modular SAT solving for IC3. In: Ray, S., Jobstmann, B. (eds.) 2013 Formal Methods in Computer-Aided Design, pp. 149–156. IEEE (2013)
4. Brown, C.W., Košta, M.: Constructing a single cell in cylindrical algebraic decomposition. J. Symb. Comput. **70**, 14–48 (2015)
5. Buchberger, B., Collins, G.E., Loos, R., Albrecht, R. (eds.): Computer Algebra. Symbolic and Algebraic Computation, Springer, Vienna (1982). https://doi.org/10.1007/978-3-7091-7551-4
6. Cavada, R., et al.: The NUXMV symbolic model checker. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 334–342. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08867-9_22
7. Caviness, B.F., Johnson, J.R. (eds.): Quantifier Elimination and Cylindrical Algebraic Decomposition. Texts and Monographs in Symbolic Computation, Springer, Vienna (2004). https://doi.org/10.1007/978-3-7091-9459-1
8. Cimatti, A., Griggio, A., Irfan, A., Roveri, M., Sebastiani, R.: Invariant checking of NRA transition systems via incremental reduction to LRA with EUF. In: Legay, A., Margaria, T. (eds.) TACAS 2017. LNCS, vol. 10205, pp. 58–75. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54577-5_4
9. Cimatti, A., Griggio, A., Sebastiani, R.: Efficient interpolant generation in satisfiability modulo theories. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 397–412. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78800-3_30
10. Collins, G.E.: Quantifier elimination for real closed fields by cylindrical algebraic decompostion. In: Brakhage, H. (ed.) GI-Fachtagung 1975. LNCS, vol. 33, pp. 134–183. Springer, Heidelberg (1975). https://doi.org/10.1007/3-540-07407-4_17

11. Craig, W.: Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. J. Symbolic Logic **22**(3), 269–285 (1957)
12. Dai, L., Xia, B., Zhan, N.: Generating non-linear interpolants by semidefinite programming. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 364–380. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39799-8_25
13. de Moura, L., Jovanović, D.: A model-constructing satisfiability calculus. In: Giacobazzi, R., Berdine, J., Mastroeni, I. (eds.) VMCAI 2013. LNCS, vol. 7737, pp. 1–12. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-35873-9_1
14. Dutertre, B.: Yices 2.2. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 737–744. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08867-9_49
15. Dutertre, B.: Solving exists/forall problems with Yices. In: 13th International Workshop on Satisfiability Modulo Theories (2015)
16. Eén, N., Sörensson, N.: An extensible SAT-solver. In: Giunchiglia, E., Tacchella, A. (eds.) SAT 2003. LNCS, vol. 2919, pp. 502–518. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24605-3_37
17. Eén, N., Sörensson, N.: Temporal induction by incremental SAT solving. Electron. Notes Theor. Comput. Sci. **89**(4), 543–560 (2003)
18. Gan, T., Dai, L., Xia, B., Zhan, N., Kapur, D., Chen, M.: Interpolant synthesis for quadratic polynomial inequalities and combination with *EUF*. In: Olivetti, N., Tiwari, A. (eds.) IJCAR 2016. LNCS (LNAI), vol. 9706, pp. 195–212. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-40229-1_14
19. Gan, T., Xia, B., Xue, B., Zhan, N., Dai, L.: Nonlinear Craig interpolant generation. In: Lahiri, S.K., Wang, C. (eds.) CAV 2020. LNCS, vol. 12224, pp. 415–438. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-53288-8_20
20. Gao, S., Zufferey, D.: Interpolants in nonlinear theories over the reals. In: Chechik, M., Raskin, J.-F. (eds.) TACAS 2016. LNCS, vol. 9636, pp. 625–641. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49674-9_41
21. Gerhold, S., Kauers, M.: A procedure for proving special function inequalities involving a discrete parameter. In: Gao, X.-S., Labahn, G. (eds.) Proceedings of the 2005 International Symposium on Symbolic and Algebraic Computation, pp. 156–162 (2005)
22. Graham-Lengrand, S., Jovanović, D., Dutertre, B.: Solving Bitvectors with MCSAT: explanations from bits and pieces. In: Peltier, N., Sofronie-Stokkermans, V. (eds.) IJCAR 2020. LNCS (LNAI), vol. 12166, pp. 103–121. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-51074-9_7
23. Henzinger, T.A., Jhala, R., Majumdar, R., McMillan, K.L.: Abstractions from proofs. ACM SIGPLAN Not. **39**(1), 232–244 (2004)
24. Hoder, K., Bjørner, N.: Generalized property directed reachability. In: Cimatti, A., Sebastiani, R. (eds.) SAT 2012. LNCS, vol. 7317, pp. 157–171. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31612-8_13
25. Hoenicke, J., Schindler, T.: Efficient interpolation for the theory of arrays. In: Galmiche, D., Schulz, S., Sebastiani, R. (eds.) IJCAR 2018. LNCS (LNAI), vol. 10900, pp. 549–565. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-94205-6_36
26. Huang, G.: Constructing Craig interpolation formulas. In: Du, D.-Z., Li, M. (eds.) COCOON 1995. LNCS, vol. 959, pp. 181–190. Springer, Heidelberg (1995). https://doi.org/10.1007/BFb0030832

27. Jovanović, D.: Solving nonlinear integer arithmetic with MCSAT. In: Bouajjani, A., Monniaux, D. (eds.) VMCAI 2017. LNCS, vol. 10145, pp. 330–346. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-52234-0_18
28. Jovanovic, D., Barrett, C., De Moura, L.: The design and implementation of the model constructing satisfiability calculus. In: Ray, S., Jobstmann, B. (eds.) 2013 Formal Methods in Computer-Aided Design, pp. 173–180. IEEE (2013)
29. Jovanović, D., de Moura, L.: Solving non-linear arithmetic. In: Gramlich, B., Miller, D., Sattler, U. (eds.) IJCAR 2012. LNCS (LNAI), vol. 7364, pp. 339–354. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31365-3_27
30. Jovanović, D., Dutertre, B.: Property-directed k-induction. In: Piskac, R., Talupur, M., Veith, H. (eds.) 2016 Formal Methods in Computer-Aided Design (FMCAD), pp. 85–92. IEEE (2016)
31. Jovanović, D., Dutertre, B.: LibPoly: a library for reasoning about polynomials. In: Proceedings 15th International Workshop on Satisfiability Modulo Theories (SMT 2017) (2017)
32. Kapur, D., Majumdar, R., Zarba, C.G.: Interpolation for data structures. In: Young, M., Devanbu, P. (eds.) Proceedings of the 14th ACM SIGSOFT International Symposium on Foundations of Software Engineering, pp. 105–116 (2006)
33. Komuravelli, A., Gurfinkel, A., Chaki, S.: SMT-based model checking for recursive programs. Formal Methods Syst. Des. **48**(3), 175–205 (2016). https://doi.org/10.1007/s10703-016-0249-4
34. Krajíček, J.: Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic. J. Symbolic Logic **62**(2), 457–486 (1997)
35. Kupferschmid, S., Becker, B.: Craig interpolation in the presence of non-linear constraints. In: Fahrenberg, U., Tripakis, S. (eds.) FORMATS 2011. LNCS, vol. 6919, pp. 240–255. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-24310-3_17
36. Mahdi, A., Scheibler, K., Neubauer, F., Fränzle, M., Becker, B.: Advancing software model checking beyond linear arithmetic theories. In: Bloem, R., Arbel, E. (eds.) HVC 2016. LNCS, vol. 10028, pp. 186–201. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-49052-6_12
37. McMillan, K.L.: An interpolating theorem prover. Theor. Comput. Sci. **345**(1), 101–121 (2005)
38. McMillan, K.L.: Quantified invariant generation using an interpolating saturation prover. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 413–427. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78800-3_31
39. McMillan, K.L.: Interpolation: proofs in the service of model checking. In: Hanbook of Model-Checking. Springer (2014)
40. Mishra, B.: Algorithmic Algebra. Springer, New York (1993). https://doi.org/10.1007/978-1-4612-4344-1
41. Pudlák, P.: Lower bounds for resolution and cutting plane proofs and monotone computations. J. Symbolic Logic **62**(3), 981–998 (1997)
42. Schindler, T., Jovanović, D.: Selfless interpolation for infinite-state model checking. In: VMCAI 2018. LNCS, vol. 10747, pp. 495–515. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-73721-8_23