

Satisfiability of Non-Linear Transcendental Arithmetic as a Certificate Search Problem

Enrico Lipparini¹[*ORCID*:0009–0009–0428–4403] and Stefan
Ratschan²[*ORCID*:0000–0003–1710–1513]

¹ DIBRIS, University of Genoa, Italy

² Institute of Computer Science of the Czech Academy of Sciences

Abstract. For typical first-order logical theories, satisfying assignments have a straightforward finite representation that can directly serve as a certificate that a given assignment satisfies the given formula. For non-linear real arithmetic with transcendental functions, however, no general finite representation of satisfying assignments is available. Hence, in this paper, we introduce a different form of satisfiability certificate for this theory, formulate the satisfiability verification problem as the problem of searching for such a certificate, and show how to perform this search in a systematic fashion. This does not only ease the independent verification of results, but also allows the systematic design of new, efficient search techniques. Computational experiments document that the resulting method is able to prove satisfiability of a substantially higher number of benchmark problems than existing methods.

1 Introduction

SAT modulo theories (SMT) is the problem of checking whether a given quantifier-free first-order formula with both propositional and theory variables is satisfiable in a specific first-order theory. In this paper, we consider the case of $\text{SMT}(\mathcal{NTA})$, non-linear real arithmetic augmented with trigonometric and exponential transcendental functions. This problem is particularly important in the verification of hybrid systems and in theorem proving. Unfortunately, \mathcal{NTA} is a very challenging theory. Indeed, it is undecidable [26], and, moreover, there is no known finite representation of satisfying assignments that could act as a direct certificate of satisfiability. This does not only make it difficult for an SMT-solver to prove satisfiability, but also raises the question of how to verify the result given by an SMT-solver.

In this paper, we introduce the notion of a satisfiability certificate for \mathcal{NTA} . Such a certificate allows independent entities to verify the satisfiability of a given input formula without having to re-do a full check of its satisfiability. More specifically, based on such a certificate, the check of satisfiability is both easier in terms of computational effort and effort needed to implement the checker and to ensure its correctness. The certificate will be based on the notion of topological

degree [11,1,12], generalizing the idea that a sign change of a continuous function f implies satisfiability of $f = 0$. The basic tool for checking correctness of the certificate is interval arithmetic [28,25,24].

The idea to verify satisfiability of $\text{SMT}(\mathcal{N}\mathcal{T}\mathcal{A})$ in such a way, is not new [21]. However, the formulation as the problem of searching for a certificate is. In addition to the possibility of independent verification, such a formulation makes the corresponding search problem explicit. This allows us to introduce new, efficient search heuristics that guide the algorithm toward finding a certificate and prevent the procedure from getting stuck in computation that later turns out to not to lead to success.

We have implemented our method in the tool UGOTNL [21] and present computational experiments with different heuristics configurations over a wide variety of $\mathcal{N}\mathcal{T}\mathcal{A}$ benchmarks. The experimental results show that this new version of UGOTNL outperforms the previous version, making it—to the best of our knowledge—the most effective solver for proving satisfiability of $\mathcal{N}\mathcal{T}\mathcal{A}$ problems.

It is possible to integrate the resulting method into a conflict-driven clause learning (CDCL) type SMT solver [21]. However, in order to keep the focus of the paper on the concern of certificate search, we ignore this possibility, here.

Content. The paper is organized as follows: In Section 2 we provide the necessary background. In Section 3 we give the formal definitions of *certifying SMT solver* and of *satisfiability certificate* in $\text{SMT}(\mathcal{N}\mathcal{T}\mathcal{A})$. In Section 4 we outline our method for searching for a certificate, and in Section 5 we illustrate the heuristics that we introduce in detail. In Section 6 we experimentally evaluate our method. In Section 7 we discuss related work. Finally, in Section 8, we draw some conclusions.

2 Preliminaries

We work in the context of *Satisfiability Modulo Theories* (SMT). Our theory of interest is the quantifier-free theory of non-linear real arithmetic augmented with trigonometric and exponential transcendental functions, $\text{SMT}(\mathcal{N}\mathcal{T}\mathcal{A})$. We assume that the reader is familiar with standard SMT terminology [5].

Notation. We denote $\text{SMT}(\mathcal{N}\mathcal{T}\mathcal{A})$ -formulas by ϕ, ψ , clauses by C_1, C_2 , literals by l_1, l_2 , real-valued variables by x_1, x_2, \dots , constants by a, b , intervals of real values by $I = [a, b]$, boxes by $B = I_1 \times \dots \times I_n$, logical terms with addition, multiplication and transcendental function symbols by f, g , and multivariate real functions with F, G, H . For any formula ϕ , we denote by $\text{vars}_{\mathcal{R}}(\phi)$ the set of its real-valued variables. When there is no risk of ambiguity we write f, g to also denote the real-valued functions corresponding to the standard interpretation of the respective terms. We assume that formulas are in Conjunctive Normal Form (CNF) and that their atoms are in the form $f \bowtie 0$, with $\bowtie \in \{=, \leq, <\}$. We remove the negation symbol by rewriting every occurrence of $\neg(f = 0)$ as $(f < 0 \vee 0 < f)$ and distributing \neg over inequalities.

Points and boxes. Since we have an order on the real-valued variables x_1, x_2, \dots , for any set of variables $V \subseteq \{x_1, x_2, \dots\}$ we can view an assignment $p : V \rightarrow \mathbb{R}$ equivalently as the $|V|$ -dimensional point $p \in \mathbb{R}^{|V|}$, and an *interval assignment* $B : V \rightarrow \{[a, b] : a, b \in \mathbb{R}\}$ equivalently as the $|V|$ -dimensional box $B \subseteq \mathbb{R}^{|V|}$. By abuse of notation, we will use both representations interchangeably, using the type \mathcal{R}^V both for assignments in $V \rightarrow \mathbb{R}$ and points in $\mathbb{R}^{|V|}$, and the type \mathcal{B}^V both for interval assignments in $V \rightarrow \{[a, b] : a, b \in \mathbb{R}\}$ and corresponding boxes. This will allow us to apply mathematical notions usually defined on points or boxes to such assignments, as well. Given a point $p \in \mathcal{R}^V$, and a subset $V' \subseteq V$, we denote by $proj_{V'}(p) \in \mathcal{R}^{V'}$ the projection of p to the variables in V' , that is, for all $v \in V'$, $proj_{V'}(p)(v) := p(v)$.

Systems of equations and inequalities. We say that a formula ϕ that contains only conjunctions of atoms in the form $f = 0$ and $g \leq 0$ is a *system of equations and inequalities*. If ϕ contains only equations (inequalities) then we say it is a *system of equations (inequalities)*. A system of equations $f_1 = 0 \wedge \dots \wedge f_n = 0$, where the f_1, \dots, f_n are terms in the variables x_1, \dots, x_m , can be seen in an equivalent way as the equation $F = 0$, where F is the real-valued function $F := f_1 \times \dots \times f_n : \mathbb{R}^m \rightarrow \mathbb{R}^n$ and 0 is a compact way to denote the point $(0, \dots, 0) \in \mathbb{R}^n$. Analogously, we can see a system of inequalities $g_1 \leq 0 \wedge \dots \wedge g_k \leq 0$ as the inequality $G \leq 0$, where G is the real-valued function $G := g_1 \times \dots \times g_k : \mathbb{R}^m \rightarrow \mathbb{R}^k$ and \leq is defined element-wise. We will write $eq(\phi)$ for the function F defined by the equations in the formula ϕ and $ineq(\phi)$ for the function G defined by the inequalities in ϕ . The handling of strict inequalities would be an easy, but technical extension of our method, which we avoid to stream-line the presentation.

Dulmage-Mendelsohn decomposition. Given a system of equations ϕ , it is possible to construct an associated bipartite graph \mathcal{G}_ϕ that represent important structural properties of the system of equations. This graph has one vertex per equation, one vertex per variable, and an edge between a variable x_i and an equation $f_j = 0$ iff x_i appears in f . The Dulmage–Mendelsohn decomposition [10,2] is a canonical decomposition from the field of matching theory that partitions the system into three parts: an over-constrained subsystem (more equalities than variables), an under-constrained subsystem (less equalities than variables), and a well-constrained subsystem (as many equalities as variables, and contains no over-constrained subsystem, i.e. it satisfies the Hall property [17]).

Example 1. Let $\phi := x - \tan(y) = 0 \wedge z^2 = 0 \wedge w = 0 \wedge \sin(w) = 0$. Through the DM-decomposition we obtain an under-constrained sub-system $x - \tan(y) = 0$ (two variables, one equation), a well-constrained sub-system $z^2 = 0$ (one variable, one equation), and an over-constrained sub-system $w = 0 \wedge \sin(w) = 0$ (one variable, two equations).

Topological degree. The notion of the degree of a continuous function (also called the topological degree) comes from differential topology [11]. For a continuous function $F : B \subseteq \mathbb{R}^n \rightarrow \mathbb{R}^n$, such that $0 \notin F(\partial B)$ (where ∂B is the topological boundary of B), the degree $\text{deg}(F, B, 0)$ is a computable [1,12] integer. This integer provides information about the roots of F in B , and can be seen as a generalization of the intermediate value theorem to higher-dimensional functions. In analogy to the fact that opposite signs of a continuous function on the endpoints of an interval imply the existence of a zero within the interval, $\text{deg}(F, B, 0) \neq 0$ implies that F has a root in B . The converse is not true, and the existence of a root does not imply nonzero degree in general. Still, if a box contains one isolated zero with non-singular Jacobian matrix, then the topological degree is non-zero [11]. For alternatives to the topological degree test see our discussion of related work.

Interval Arithmetic. The basic algorithmic tool that underlies our approach is floating point interval arithmetic (\mathcal{IA} [28,25,24] which, given a box B and an \mathcal{NTA} -term representing a function H , is able to compute an interval $\mathcal{IA}_H(B)$ that over-approximates the range $\{H(x) \mid x \in B\}$ of H over B . Since this is based on floating point arithmetic, the time needed for computing $\mathcal{IA}_H(B)$ does not grow with the size of the involved numbers. Moreover conservative rounding guarantees correctness under the presence of round-off errors. In the paper, we will use interval arithmetic within topological degree computation [12], and as a tool to prove the validity of inequalities on boxes.

Robustness. We say that a formula ϕ is robust if there exists some $\epsilon > 0$ such that ϕ is satisfiable iff every ϵ -perturbation of ϕ is satisfiable (for the precise definition of ϵ -perturbation see [13]). If ϕ is both robust and (un)satisfiable, we say that it is robustly (un)sat.

Relation between robustness and system of equations: An over-constrained system of equations is never robustly sat [13, Lemma 5]. It easily follows that a system of equations that contains an over-constrained sub-system (in the sense of the Dulmage-Mendelsohn decomposition) is never robustly sat as well.

Relation between robustness and topological degree: Even in the case of an isolated zero, the test for non-zero topological degree can fail if the system is non-robust. For example, the function $F(x) \equiv x^2$ has topological degree 0 in the interval $[-1, 1]$, although the equality $x^2 = 0$ has an isolated zero in this interval. However, the zero of $x^2 = 0$ is not robust: it can vanish under arbitrarily small changes of the function denoted by the left-hand side x^2 . It can be shown that the topological degree test is able to prove satisfiability in all robust cases for a natural formalization of the notion of robustness [13]. We will not provide such a formalization, here, but use robustness as an intuitive measure for the potential success when searching for a certificate.

Logic-To-Optimization. While symbolic methods usually struggle dealing with $\mathcal{N}\mathcal{T}\mathcal{A}$, numerical methods, albeit inexact, can handle transcendental functions efficiently. For this reason, an SMT solver can benefit from leveraging numerical techniques. In the Logic-To-Optimization approach [21,15], an $\text{SMT}(\mathcal{N}\mathcal{T}\mathcal{A})$ -formula ϕ in m variables is translated into a real-valued non-negative function $\mathcal{L}2\mathcal{O}(\phi) \equiv H : \mathbb{R}^m \mapsto \mathbb{R}^{\geq 0}$ such that—up to a simple translation between Boolean and real values for Boolean variables—each model of ϕ is a zero of H (but not vice-versa). When solving a satisfiability problem, one can try to first numerically minimize this function, and then use the obtained numerical (approximate) solution to prove, through exact symbolic methods, that the logical formula has indeed a model. For the precise definition of the operator $\mathcal{L}2\mathcal{O}$ see [21, Section 3].

3 Goal

Consider an SMT solver that takes as input some formula ϕ and as output an element of $\{\text{sat}, \text{unknown}, \text{unsat}\}$. How can we gain trust in the correctness of the result of such an SMT solver? One approach would be to ensure that the algorithm itself is correct. Another option is to provide a second algorithm whose output we compare with the original one. Both approaches are, however, very costly, and moreover, the latter approach still may be quite unreliable.

Instead, roughly following McConnell et. al. [23] (see also Figure 1), we require our solver to return—in addition to its result—some information that makes an independent check of this result easy:

Definition 1. *An SMT solver is certifying iff for an input formula ϕ , in addition to an element $r \in \{\text{sat}, \text{unknown}, \text{unsat}\}$, it returns an object w (a certificate) such that*

- (ϕ, r, w) satisfies a property W where $W(\phi, r, w)$ implies that r is a correct result for ϕ , and
- there is an algorithm (a certificate checker) that
 - takes as input a triple (ϕ, r, w) and returns \top iff $W(\phi, r, w)$, and that
 - is simpler than the SMT solver itself.




Fig. 1. Certifying SMT Solver

So, for a given formula ϕ , one can ensure correctness of the result (r, w) of a certifying SMT solver by using a certificate checker to check the property $W(\phi, r, w)$. Since the certificate checker is simpler than the SMT solver itself, the correctness check is simpler than the computation of the result itself.

The definition leaves it open, what precisely is meant by “simpler”. In general, it could either refer to the run-time of the checker, or to the effort needed for implementing the certificate checker and ensuring its correctness. The former approach is taken in computational complexity theory, the latter in contexts where correctness is the main concern [23]. Indeed, we will later see that our approach succeeds in satisfying both requirements, although we will not use complexity-theoretic measures of run-time, but will measure run-time experimentally.

The use of such certificates is ongoing research in the unsatisfiable case [4]. In the satisfiable case, for most theories, one can simply use satisfying assignments (i.e., witnesses) as certificates. Here the property W simply is the property that the given assignment satisfies the formula, which can be checked easily.

For $\text{SMT}(\mathcal{N}\mathcal{T}\mathcal{A})$, however, the situation is different: Here, no general finite representation of satisfying assignments is available. Hence one needs to use certificates of a different form. We introduce the following definition:

Definition 2. *Let ϕ be a formula in $\mathcal{N}\mathcal{T}\mathcal{A}$. A (satisfiability) certificate for ϕ is a triple (σ, ν, β) such that $W(\phi, \text{sat}, (\sigma, \nu, \beta))$ iff*

- σ is a function selecting a literal from every clause of ϕ
- ν is a variable assignment in \mathcal{R}^V assigning floating point numbers to a subset $V \subseteq \text{vars}_{\mathcal{R}}(\sigma(\phi))$ (where $\sigma(\phi)$ is a compact way of writing $\bigwedge_{C \in \phi} \sigma(C)$), s.t. $\sigma(\phi)$ contains as many equations as real-valued variables not in V .
- β is a finite set of interval assignments in $\mathcal{B}^{\text{vars}_{\mathcal{R}}(\phi) \setminus V}$ such that their set-theoretic union as boxes is again a box B_β and, for the system of equations $F := \text{eq}(\nu(\sigma(\phi)))$ and the system of inequalities $G := \text{ineq}(\nu(\sigma(\phi)))$, it holds that:
 - $0 \notin F(\partial B_\beta)$,
 - $\text{deg}(F, B_\beta, 0) \neq 0$, and
 - for every $B \in \beta$, $\mathcal{I}A_G(B) \leq 0$.

Example 2. Consider the formula

$$\begin{aligned} \phi &:= C_1 \wedge C_2 \wedge C_3 \wedge C_4 \\ C_1 &\equiv \cos(y) = 0 \vee \sin(y) = e^x & C_3 &\equiv x - y \leq \cos(z) \\ C_2 &\equiv \sin(y) = 0 \vee \cos(y) = \sin(8x^2 - z) & C_4 &\equiv x + y \geq \sin(z) \end{aligned}$$

The following (σ, ν, β) is a certificate:

- $\sigma := \{C_1 \mapsto \sin(y) = e^x ; C_2 \mapsto \cos(y) = \sin(8x^2 - z) ; C_3 \mapsto C_3 ; C_4 \mapsto C_4\}$
- $\nu := \{z \mapsto 0.2\}$
- $\beta := \{B\}$, where $B := \{x \mapsto [-0.1, 0.05] ; y \mapsto [1.4, 1.9]\}$




Fig. 2. Solution Sets of Equalities of Example Certificate

As can be seen in Figure 2, the solution sets of C_1 and C_2 cross at a unique point in B , which reflects the fact that the degree of the function $(x, y) \rightarrow (\sin(y) - e^x, \cos(y) - \sin(8x^2 - 0.2))$ is non-zero. Moreover, the inequalities C_3 and C_4 hold on all elements of the box.

Due to the properties of the topological degree and of interval arithmetic discussed in the preliminaries, we have:

Property 1. $W(\phi, \text{sat}, (\sigma, \nu, \beta))$ implies that ϕ is satisfiable.

Moreover, the topological degree can be computed algorithmically [1,12], and one can easily write a certificate checker based on such an algorithm. Hence such a triple can be used as a certificate for satisfiability.

In this paper, we will show that in addition to the discussed benefits for correctness, formulating satisfiability checking as the problem of search for such certificates also is beneficial for efficiency of the SMT solver itself. Since we will concentrate on satisfiability, we will simply ignore the case when an SMT solver returns `unsat`, so the reader can simply assume that an SMT solver such as the one from Figure 1 only returns an element from the set $\{\text{sat}, \text{unknown}\}$.

4 Method

Our goal is to find a triple (σ, ν, β) that is a certificate of satisfiability for a given formula ϕ . So we have a search problem. In order to make this search as efficient as possible, we want to guide the search toward a triple that indeed turns out to be a certificate, and for which the corresponding conditions are computationally easy to check.

Intuitively, we view the search for a certificate as a hierarchy of nested search problems, where the levels of this hierarchy correspond to the individual components of certificates. We formalize this using a search tree whose nodes on the i -th level are labeled with i -tuples containing the first i elements of the tuple searched for, starting with the root node that is labeled with the empty tuple $()$. The tree

will be spanned by a function ch that assigns to each node (c_1, \dots, c_i) of the tree a sequence $\langle x_1, \dots, x_n \rangle$ of possible choices for the next tuple component. Hence the children of (c_1, \dots, c_i) in the tree are $(c_1, \dots, c_i, x_1), \dots, (c_1, \dots, c_i, x_n)$. We will do depth-first search in the resulting tree, searching for a leaf labeled by a certificate of satisfiability for the input formula ϕ .

Based on the observation that on each level of the tree one has the first i components of the tuple available for determining a good sequence of choices, we will add additional information as the first tuple component in the form of a variable assignment p that satisfies the formula ϕ approximately. Hence we search for a 4-tuple (p, σ, ν, β) .

It is easy to see that it would be possible to generalize such a search tree to a more fine-grained one, where the individual levels are formed by parts of the choices described above, and where the order of those levels can be arbitrary. For example, it would be possible to first choose an interval for a variable (i.e., part of the box β), then select a literal from a certain clause (i.e., part of the selection function σ), and so on. However, in this paper, we keep these levels separated, as discussed above, in order to achieve a clear separation of concerns when exploring design choices at the individual levels.

5 Certificate Search

In this section, we will discuss possibilities for search strategies by defining for every search tree node labeled with tuple τ , the ordered sequence $ch(\tau)$ of choices for the next tuple element. Our framework allows for many more possibilities from which we choose strategies that both demonstrate the flexibility of the framework, and allow for efficient search, as will be demonstrated by the computational experiments in Section 6.

In order to be able to refer to different variants of the search strategy in the description of computational experiments, we will introduce keywords for those variants that we will write in teletype font.

5.1 Points

The points $ch() = \langle p_1, \dots, p_k \rangle$ determining the first level of the search tree are generated by an optimization problem defined on the formula ϕ following the Logic-To-Optimization approach [21]. Here we translate the satisfiability problem into a numerical minimization problem, mapping the logic formula ϕ into the non-negative real-valued function $\mathcal{L}\mathcal{O}(\phi) \equiv H : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ (called the *objective function*) such that for every satisfying assignment, this objective function is zero, and for assignments that do not satisfy the formula, the objective function is typically (but not always) non-zero. Then we find local minima of H through an unconstrained optimization algorithm such as basin hopping [30]. In our implementation, we compute $k = 100$ local minima, and process them in the order of their value.

5.2 Literals

Given a point p , we choose literal selector functions $ch(p) = \langle \sigma_1, \dots, \sigma_k \rangle$ by restricting ourselves, for each clause C , to the literals l for which the objective function restricted to l and evaluated in the point p is below a certain threshold. That is, we determine the set of approximately satisfiable literals

$$L_C := \{l \in C \mid \mathcal{L}2\mathcal{O}(l)(p) \leq \epsilon\}.$$

Our literal selector functions will then correspond to the set of all approximately satisfiable combinations,

$$\{\sigma \mid \text{for all } C \in \phi, \sigma(C) \in L_C\},$$

that is, each σ selects exactly one approximately satisfiable literal from each clause. In order to maximize the chances of choosing a better literal combination, we can sort each L_C according to the value of the respective objective functions and then choose literal combinations using the corresponding lexicographic order (we will refer to this heuristic as **(sort-literals)**).

While the point p is usually a good candidate in terms of *distance from a zero*, it can sometimes lead to an inconsistent problem:

Example 3. Consider the formula

$$\phi := C_1 \wedge C_2$$

$$C_1 \equiv (x + y = 0) \vee (x = e^{10^6 * y}) \quad C_2 \equiv (x + y \geq \epsilon_1) \vee (x = \tan(y + \epsilon_1))$$

The numerical optimizer will be tempted to return first some point p_1 such as $\{x \mapsto 1; y \mapsto -1\}$, that *almost* satisfies $(x + y = 0) \wedge (x + y \geq \epsilon_1)$, instead of a harder approximate solution involving transcendental functions and heavy approximations, such as $(x = e^{10^6 * y}) \wedge (x = \tan(y + \epsilon_1))$, that is exactly satisfiable in a point p_2 near $(0, -\pi)$.

Such inconsistencies may occur in many combinations of literals. We use a strategy that detects them in situations where for certain clauses C , the set L_C contains only one literal l . We will call such a literal l a *forced literal*, since, for every literal selector function σ , $\sigma(\phi)$ will include l . Before starting to tackle every approximately satisfiable literal combination, we first analyze the set of forced literals. We do symbolic simplifications (such as rewriting and Gaussian elimination) to check whether the set has inconsistencies that can be found at a symbolic level (as in the previous example). If the symbolic simplifications detect that the forced literals are inconsistent then we set $ch(p)$ to the empty sequence $\langle \rangle$ which causes backtracking in depth-first search. We refer to the variant of the algorithm using this check as **(check-forced-literals)**.

Filtering out over-constrained systems. Given a literal selector function σ , we analyze the structure of the system of equations formed by the equations selected by σ through the Dulmage-Mendelsohn decomposition, that uniquely

decomposes the system into a well-constrained subsystem, an over-constrained subsystem and an under-constrained subsystem. We filter out every literal combination having a non-empty over-constrained subsystem, since this leads to a non-robust sub-problem, referring to this heuristic as **(filter-overconstr)**.

5.3 Instantiations

We define the instantiations $ch(p, \sigma) = \langle \nu_1, \dots, \nu_k \rangle$ based on a sequence of sets of variables V_1, \dots, V_k to instantiate, and define $\nu_i := proj_{V_i}(p)$. The uninstantiated part of p after projection to a set of variables V_i is then $proj_{vars_{\mathcal{R}}(\phi) \setminus V_i}(p)$, which we will denote by p_{-V_i} .

For searching for the variables to instantiate, we use the Dulmage–Mendelsohn decomposition constructed in the previous level of the hierarchy. We do not want to instantiate variables appearing in the well-constrained sub-system, since doing so would make the resulting system after the instantiation over-constrained. Hence the variables to be instantiated should be chosen only from the variables occurring in the under-constrained subsystem. This substantially reduces the number of variable combinations that we can try. Denoting the variables satisfying this criterion by V_{under} , this restricts $V_i \subseteq V_{under}$, for all $i \in \{1, \dots, k\}$. This does not yet guarantee that every chosen variable combination leads to a well-constrained system after the instantiation. For example, the under-determined system of equations $x + y = 0 \wedge z + w = 0$ has four variables and two equations, but becomes over-constrained after instantiating either the two variables x and y , or the variables z and w . So, for each V_i , we further check whether the system obtained after the instantiation is well-constrained (we refer to this heuristic as **(filter-overconstr-V)**).

The method described in the previous paragraph only uses information about which equations in the system contain which variables (i.e., it deals only with the *structure* of the system, not with its *content*). Indeed, it ignores the point p .

To extract more information, we use the fact that a non-singular Jacobian matrix of a function at one of its zeros implies a non-zero topological degree wrt. every box containing this single zero [11]. So we compute a floating point approximation of the Jacobian matrix at point p (note that, in general, this matrix is non-square). Our goal is to find a set of variables V to instantiate such that the Jacobian matrix corresponding to the resulting square system at the point p_{-V} has full rank. This matrix is the square sub-matrix of the original Jacobian matrix that is the result of removing the instantiated columns.

A straight-forward way of applying the Jacobian criterion is, given random variable instantiations, to filter out instantiations whose corresponding Jacobian matrix is rank-deficient (**filter-rank-deficient**), similarly to what is done in the previous paragraph with the overconstrained filter. Note that, as the Jacobian matrix of non-well-constrained system of equations is always rank-deficient, this filter is stronger than the previous one. However, it may filter out variable instantiations that result in a non-zero degree (e.g., the function x^3 has non-zero degree in $[-1, 1]$, but its Jacobian matrix at the origin is rank deficient since $f'(0) = 0$).

We can further use the information given by the Jacobian matrix not only to filter out bad variable instantiations, but also to maximize the chance of choosing good variable instantiations from the beginning. Indeed, not all variable instantiations will be equally promising, and it makes sense to head for an instantiation such that the resulting square matrix not only has full rank, but—in addition—is far from being rank-deficient (i.e., it is as robust as possible). We can do so by modifying Kearfott’s method [19, Method 2] which fixes the coordinates most tangential to the orthogonal hyper-plane of F in p by first computing an approximate basis of the null space of the Jacobian matrix in the point, and then choosing the variables with maximal sum of absolute size. We use a modification of the method that uses a variable ordering w.r.t. this sum, and then extracts the sets of variables V_1, V_2, \dots in decreasing order w.r.t. the cumulative sum of the value of the variables in each set. We refer to this heuristic as (**Kearfott-ordering**).

5.4 Box

We construct boxes around p_{-V} , where V is the set of variables ν instantiates, that is, $\nu \in \mathcal{R}^V$. So we define $ch(p, \sigma, \nu) := \langle \beta_1, \dots, \beta_k \rangle$ s.t. for all $i \in \{1, \dots, k\}$, for all $B \in \beta_i$, $B \in \mathcal{B}^{vars\kappa(\phi)-V}$ and $p_{-V} \in \bigcup_{B \in \beta_i} B$.

We use two different methods, (**eps-inflation**) and (**box-gridding**):

- Epsilon-inflation [22] is a method to construct incrementally larger boxes around a point. In this case, the β_1, \dots, β_k will each just contain one single box B_i defined as the box centered at p_{-V} having side length $2^i \epsilon$, where, in our setting, $\epsilon = 10^{-20}$. We terminate the iteration if either $\mathcal{I}A_G(B_i) \leq 0$ and $deg(F, B_i, 0) \neq 0$, in which case we found a certificate, or we reach an iteration limit (in our setting when $2^i \epsilon > 1$).
- Box-gridding is a well-known technique from the field of interval arithmetic based on iteratively refining a starting box into smaller sub-boxes. Here we use a specific version, first proposed in [13] and then implemented with some changes in [21]. In the following we roughly outline the idea behind the algorithm, and refer to the other two papers for details. We start with a grid that initially contains a starting box. We then iteratively refine the grid by splitting the starting box into smaller sub-boxes. At each step, for each sub-box B we first check whether interval arithmetic can prove that the inequalities or the equations are unsatisfiable, and, if so, we remove B from the grid. We check also whether $deg(F, B, 0) \neq 0$ and interval arithmetic can prove the satisfiability of the inequalities, and, if so, then we terminate our search, finding a certificate with the singleton $\beta_i = \{B\}$. In some cases, in order to verify the satisfiability of the inequalities, we will have to further split the box B into sub-boxes, using the set of resulting sub-boxes instead of the singleton $\{B\}$. After each step, if there are sub-boxes left in the grid, we continue the refinement process. Otherwise, if the grid is empty, we conclude that there cannot be solutions in the starting box. If a certain limit to the grid size is exceeded, we also stop the box gridding procedure without success.

For both methods, if the method stops without success, we have arrived at the last element of the sequence of choices $\langle \beta_1, \dots, \beta_k \rangle$ without finding a certificate, which results in backtracking of the depth-first search for a certificate.

Both mentioned methods have their advantages, and can be seen as complementary. Epsilon-inflation is quite fast, and performs particularly well if the solution is isolated and is near the center. However, if there are multiple solutions in a box, the topological degree test can potentially fail to detect them³, and if the solution is far from the center then we need a bigger box to encompass it, which is less likely to be successful than a smaller box, as we require the inequalities to hold everywhere in the box, and, moreover, the chance of encompassing other solutions (thus incurring in the previous problem) grows.

The box-gridding procedure, on the other side, can be quite slow, as in the worst case the number of sub-boxes explodes exponentially. However, grid refinement leads to a very accurate box search, which allows us to avoid the issues faced with epsilon inflation (i.e. multiple solutions, or a solution far from the center). Moreover, if the problem is robust, we have the theoretical guarantee that the procedure will eventually converge to a solution [13], although this does not hold in practice due to the introduced stopping criterion.

Indeed, a third approach is to combine the two methods: first use epsilon inflation, that is often able to quickly find a successful box, and, if it fails, then use the more accurate box-gridding procedure.

6 Computational Experiments

Implementation. We implemented the different heuristics presented in the paper in a prototype tool called UGOTNL (firstly presented in [21]). In order to make the results comparable with the ones obtained earlier, in addition to the search method discussed in Section 5, we preserve the following heuristics used by UGOTNL: If the local minimizer cannot find any minimum of $\mathcal{L}2\mathcal{O}(\phi)$ for which for every clause $C \in \phi$, the set of approximately satisfiable literals L_C is non-empty, we restart the procedure on every conjunction resulting from the DNF of ϕ . The tool handles strict inequalities of the form $f < 0$ directly until the box construction phase, where they are replaced by $f \leq -\varepsilon$ (with $\varepsilon = 10^{-20}$). For computing the topological degree, we use TOPDEG⁴. For the symbolic simplifications used in (`check-forced-literals`), we use the *simplify* and the *solve-eqs* tactics provided by z3 [9]⁵. For the computation of the rank used in (`filter-rank-deficient`), we observe that the rank of a matrix is equal to the number of non-zero singular values, hence we consider a matrix far from rank-deficiency iff all its singular values are bigger than some threshold (to account for approximation errors). We use a threshold widely used by algorithms for

³ For example, for $f(x) = x^2 - 1$, $\text{deg}(f, [-10, 10], 0) = 0$, while $\text{deg}(f, [-10, 0], 0) = -1$, and $\text{deg}(f, [0, 10], 0) = 1$.

⁴ Available at <https://www.cs.cas.cz/~ratschan/topdeg/topdeg.html>.

⁵ For a description of the two tactics: <https://microsoft.github.io/z3guide/docs/strategies/summary>. The version of z3 used is 4.5.1.0.

N. solved	Heuristics			(id.)
	Literals	Instantiations	Boxes	
323			(box-gridding)	(1.a.)
355			(eps-inflation)	(1.b.)
356			(eps-inflation) (box-gridding)	(1.c.)
362	(sort-literals)		(eps-inflation)	(2.b.)
361	(sort-literals)		(eps-inflation) (box-gridding)	(2.c.)
370	(sort-literals) (filter-overconstr)		(eps-inflation)	(3.b.)
367	(sort-literals) (filter-overconstr)		(eps-inflation) (box-gridding)	(3.c.)
406	(sort-literals) (filter-overconstr) (check-forced-literals)		(eps-inflation)	(4.b.)
410	(sort-literals) (filter-overconstr) (check-forced-literals)		(eps-inflation) (box-gridding)	(4.c.)
409	(sort-literals) (filter-overconstr) (check-forced-literals)	(Kearfott-ordering)	(eps-inflation)	(5.b.)
412	(sort-literals) (filter-overconstr) (check-forced-literals)	(Kearfott-ordering)	(eps-inflation) (box-gridding)	(5.c.)
424	(sort-literals) (filter-overconstr) (check-forced-literals)	(Kearfott-ordering) (filter-overconstr-V)	(eps-inflation)	(6.b.)
426	(sort-literals) (filter-overconstr) (check-forced-literals)	(Kearfott-ordering) (filter-overconstr-V)	(eps-inflation) (box-gridding)	(6.c.)
427	(sort-literals) (filter-overconstr) (check-forced-literals)	(Kearfott-ordering) (filter-overconstr-V) (filter-rank-deficient)	(eps-inflation)	(7.b.)
426	(sort-literals) (filter-overconstr) (check-forced-literals)	(Kearfott-ordering) (filter-overconstr-V) (filter-rank-deficient)	(eps-inflation) (box-gridding)	(7.c.)
441	Virtual best			

Fig. 3. Summary of the results for different heuristics configurations. Each row corresponds to a configuration. The first column from the left contains the number of benchmarks solved; the central columns indicate the heuristics used, separated by search level; the last column contains an identifier of the configuration. The last row is for the virtual best of the different configurations.

determining the matrix rank, which is $\sigma_{\max} \dim(A) \varepsilon$, where σ_{\max} is the largest singular value of A , and ε is the machine epsilon.

Setup. We run the experiments⁶ on a cluster of identical machines equipped with 2.6GHz AMD Opteron 6238 processors. We set a time limit of 1000 seconds, and a memory limit of 2Gb. We considered all SMT($\mathcal{N}\mathcal{T}\mathcal{A}$) benchmarks from the dReal distribution[16] and other SMT($\mathcal{N}\mathcal{T}\mathcal{A}$) benchmarks coming from the discretization of Bounded Model Checking of hybrid automata [27,3], totaling 1931 benchmarks. All of these benchmarks come with "unknown" status. According to experiments performed on other solvers (CVC5, MATHSAT, DREAL), among these benchmarks 736 (respectively, 136) are claimed to be unsatisfiable (satisfiable) by at least one solver⁷. We tested our tool with different heuristics configurations (Figure 3), and, for each configuration, we checked that our tool never contradict the other tools. We have arranged the heuristics into 3 columns (Literals, Instantiations, and Boxes) according to the search level they are used in. As the number of possible configurations is quite high, we proceed as follows: We start with the simpler configurations (just one method for finding a box that contains a solution), and then we add heuristics.

Results. In the first configurations we tested the 3 possible ways to search for a box. We note that (**box-gridding**) (1.a.) performs considerably worse than the other two, (**eps-inflation**) (1.b.) and (**eps-inflation**) + (**box-gridding**) (1.c.), which produce comparable results. Because of that, and for readability's sake, we did not use (**box-gridding**) alone with other heuristics in the next configurations, but only considered the other two options. We then added heuristics based on the following criteria: first heuristics for the "Literals" choice, then heuristics for the "Instantiations" choice, and first ordering heuristics (i.e. (**sort-literals**) and (**Kearfott-ordering**)), then filtering heuristics (all the others). At every new heuristic added, we see that the number of benchmarks solved grows regardless of the "Boxes" choice, with the best configuration reaching 427 benchmarks using 7 heuristics. If we consider the virtual best (i.e. run in parallel all the configurations and stop as soon as a certificate is found) we are able to solve 441 benchmarks. This is because in cases such as (**eps-inflation**) vs. (**eps-inflation**) + (**box-gridding**), or such as (**filter-overconstr-V**) vs. (**filter-rank-deficient**), there is no dominant choice, with each configuration solving benchmarks that the other does not solve and vice-versa.

Discussion. The first configuration (1.a.) essentially uses a method proposed earlier [21] and implemented in a tool called UGOTNL_{EAGER} (of which the tool presented in this paper is an upgrade). Already in the previous paper, UGOTNL_{EAGER} outperformed the other solvers able to prove satisfiability in SMT($\mathcal{N}\mathcal{T}\mathcal{A}$), solving more than three times the benchmarks than MATHSAT [8], CVC5 [20], and ISAT3 [14], and almost as twice as the benchmarks solved by the *lazy* version MATHSAT+UGOTNL (where UGOTNL had been integrated *lazily* inside MATHSAT). Here we show that the new heuristics introduced further improve

⁶ The results of the experiments are available at <https://doi.org/10.5281/zenodo.7774117>

⁷ For the results of such experiments, see [21].

the performances of our tool, that is now able to solve around 100 benchmarks more.

Run-time of the certificate checker. In Section 3 we claimed that, with our approach, checking a certificate requires less run-time than the certificate search itself. Here we experimentally quantify this amount: for each benchmark solved by the best configuration (7.b.), we observe the run-time required to check the certificate (which amounts, essentially, to the computation of topological degree and interval arithmetic for the successful box). In terms of median (respectively, mean), checking the certificate requires 0.10% (1.07%) of the run-time used by the solver.

7 Related Work

One strategy for proving satisfiability in $\text{SMT}(\mathcal{NTA})$ is to prove a stricter requirement that implies satisfiability, but is easier to check. For example, one can prove that *all* elements of a set of variable assignments satisfy the given formula [14], or that a given variable assignment satisfies the formula for *all possible interpretations* of the involved transcendental functions within some bounds [7]. Such methods may be quite efficient in proving satisfiability of formulas with inequalities only, since those often have full-dimensional solution sets. However, such methods usually fail to prove satisfiability of equalities, except for special cases with straightforward rational solutions.

Computation of formally verified solutions of square systems of equations is a classical topic in the area of interval analysis [28,25,24]. Such methods usually reduce the problem either to fixpoint theorems such as Brouwer’s fixpoint theorem or special cases of the topological degree, for example, Miranda’s theorem. Such tests are easier to implement, but less powerful than the topological degree (the former fails to verify equalities with double roots, such as $x^3 = 0$, and the latter requires the solution sets of the individual equalities to roughly lie normal to the axes of the coordinate system).

In the area of rigorous global optimization, such techniques are applied [18,19] to conjunctions of equalities and inequalities in a similar way as in this paper, but with a slightly different goal: to compute rigorous upper bounds on the global minimum of an optimization problem. This minimum is often attained at the boundary of the solution set of the given inequalities, whereas satisfiability is typically easier to prove far away from this boundary.

We are only aware of two approaches that extend verification techniques for square systems of equations to proving satisfiability of quantifier-free non-linear arithmetic [29,21], one [29] being restricted to the polynomial case, and the other one also being able to handle transcendental function symbols. Neither approach is formulated in the form of certificate search. However, both could be interpreted as such, and both could be extended to return a certificate. The present paper actually does this for the second approach [21], and demonstrates that this does not only ease the independent verification of results, but also allows

the systematic design of search techniques that result in significant efficiency improvements.

An alternative approach is to relax the notation of satisfiability, for example using the notion of δ -satisfiability [16,6], that does *not* guarantee that the given formula is satisfiable, but only that the formula is not too far away from a satisfiable one, for a suitable formalization of the notion of “not too far away”. Another strategy is to return candidate solutions in the form of bounds that guarantee that certain efforts to prove unsatisfiability within those bounds fail [14].

8 Conclusions

We introduced a form of satisfiability certificate for $\text{SMT}(\mathcal{N}\mathcal{T}\mathcal{A})$ and formulated the satisfiability verification problem as the problem of searching for such a certificate. We showed how to perform this search in a systematic fashion introducing new and efficient search techniques. Computational experiments document that the resulting method is able to prove satisfiability of a substantially higher number of benchmark problems than existing methods.

Acknowledgments

The authors thank Alessandro Cimatti, Alberto Griggio, and Roberto Sebastiani for helpful discussions on the topic of the paper. The work of Stefan Ratschan was supported by the project GA21-09458S of the Czech Science Foundation GA ĀR and institutional support RVO:67985807.

References

1. Oliver Aberth. Computation of topological degree using interval arithmetic, and applications. *Math. Comput.*, 62(205):171–178, jan 1994.
2. Samy Ait-Aoudia, Roland Jégou, and Dominique Michelucci. Reduction of constraint systems. *CoRR*, abs/1405.6131, 2014.
3. Stanley Bak, Sergiy Bogomolov, and Taylor T. Johnson. Hyst: A source transformation and translation tool for hybrid automaton models. In *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control, HSCC '15*, page 128–133, New York, NY, USA, 2015. Association for Computing Machinery.
4. Haniel Barbosa, Andrew Reynolds, Gereon Kremer, Hanna Lachnitt, Aina Niemetz, Andres Nötzli, Alex Ozdemir, Mathias Preiner, Arjun Viswanathan, Scott Viteri, Yoni Zohar, Cesare Tinelli, and Clark Barrett. Flexible proof production in an industrial-strength smt solver. In *Automated Reasoning: 11th International Joint Conference, IJCAR 2022, Haifa, Israel, August 8–10, 2022, Proceedings*, page 15–35, Berlin, Heidelberg, 2022. Springer-Verlag.
5. Clark W. Barrett, Roberto Sebastiani, Sanjit A. Seshia, and Cesare Tinelli. Satisfiability modulo theories. In *Handbook of Satisfiability*, 2021.

6. Franz Brauße, Konstantin Korovin, Margarita Korovina, and Norbert Müller. *The ksmt Calculus Is a δ -complete Decision Procedure for Non-linear Constraints*, volume 12699 of *Lecture Notes in Computer Science*, pages 113–130. Springer, 2021.
7. Alessandro Cimatti, Alberto Griggio, Ahmed Irfan, Marco Roveri, and Roberto Sebastiani. Incremental linearization for satisfiability and verification modulo non-linear arithmetic and transcendental functions. *ACM Trans. Comput. Logic*, 19(3), aug 2018.
8. 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.
9. 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.
10. A. L. Dulmage and N. S. Mendelsohn. Coverings of bipartite graphs. *Canadian Journal of Mathematics*, 10:517–534, 1958.
11. Irene Fonseca and Wilfrid Gangbo. *Degree Theory in Analysis and Applications*. Clarendon Press, Oxford, 1995.
12. Peter Franek and Stefan Ratschan. Effective topological degree computation based on interval arithmetic. *Mathematics of Computation*, 84:1265–1290, 2015.
13. Peter Franek, Stefan Ratschan, and Piotr Zgliczynski. Quasi-decidability of a fragment of the first-order theory of real numbers. *Journal of Automated Reasoning*, 57(2):157–185, 2016.
14. Martin Fränzle, Christian Herde, Tino Teige, Stefan Ratschan, and Tobias Schubert. Efficient solving of large non-linear arithmetic constraint systems with complex boolean structure. *JSAT*, 1:209–236, 05 2007.
15. Zhoulai Fu and Zhendong Su. Xsat: A fast floating-point satisfiability solver. In *CAV (2)*, volume 9780 of *Lecture Notes in Computer Science*, pages 187–209. Springer, 2016.
16. Sicun Gao, Soonho Kong, and Edmund M. Clarke. dReal: An SMT solver for nonlinear theories over the reals. In Maria Paola Bonacina, editor, *Automated Deduction—CADE-24*, pages 208–214, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
17. P. Hall. On representatives of subsets. *Journal of the London Mathematical Society*, s1-10(1):26–30, 1935.
18. E. Hansen. *Global Optimization Using Interval Analysis*. Marcel Dekker, New York, 1992.
19. R. Baker Kearfott. On proving existence of feasible points in equality constrained optimization problems. *Mathematical Programming*, 83(1):89–100, 1998.
20. Gereon Kremer, Andrew Reynolds, Clark Barrett, and Cesare Tinelli. Cooperating techniques for solving nonlinear real arithmetic in the cvc5 SMT solver (system description). In Jasmin Blanchette, Laura Kovács, and Dirk Pattinson, editors, *Automated Reasoning*, pages 95–105, Cham, 2022. Springer International Publishing.
21. 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.
22. G. Mayer. Epsilon-inflation in verification algorithms. *Journal of Computational and Applied Mathematics*, 60:147–169, 1994.

23. R.M. McConnell, K. Mehlhorn, S. Näher, and P. Schweitzer. Certifying algorithms. *Computer Science Review*, 5(2):119 – 161, 2011.
24. Ramon E. Moore, R. Baker Kearfott, and Michael J. Cloud. *Introduction to Interval Analysis*. SIAM, 2009.
25. Arnold Neumaier. *Interval Methods for Systems of Equations*. Cambridge University Press, Cambridge, 1991.
26. Daniel Richardson. Some undecidable problems involving elementary functions of a real variable. *J. Symb. Log.*, 33(4):514–520, 1968.
27. Nima Roohi, Pavithra Prabhakar, and Mahesh Viswanathan. Hare: A hybrid abstraction refinement engine for verifying non-linear hybrid automata. In Axel Legay and Tiziana Margaria, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 573–588, Berlin, Heidelberg, 2017. Springer Berlin Heidelberg.
28. Siegfried M. Rump. Verification methods: Rigorous results using floating-point arithmetic. *Acta Numerica*, page 3–4, 2010.
29. Tung Vu Xuan, To Khanh, and Mizuhito Ogawa. raSAT: an SMT solver for polynomial constraints. *Formal Methods in System Design*, 51, 12 2017.
30. David J. Wales and Jonathan P. K. Doye. Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms. *The Journal of Physical Chemistry A*, 101(28):5111–5116, 1997.