

Exploiting Strict Constraints in the Cylindrical Algebraic Covering

Philipp Bär¹, Jasper Nalbach^{1,*}, Erika Ábrahám¹ and Christopher W. Brown²

¹RWTH Aachen University, Germany

²United States Naval Academy, USA

Abstract

One of the few available complete methods for checking the satisfiability of sets of polynomial constraints over the reals is the *cylindrical algebraic covering (CALC)* method. In this paper, we propose an extension for this method to exploit the *strictness* of input constraints for reducing the computational effort. We illustrate the concepts on a multidimensional example and provide experimental results to evaluate the usefulness of our proposed extension.

Keywords

Satisfiability checking, SMT solving, Real algebra, Cylindrical algebraic covering, Strict constraints

1. Introduction

Quantifier-free (non-linear) real-algebraic (QFNRA) formulas are Boolean combinations of polynomial constraints over the reals. Efficiently determining the *satisfiability* of such formulas has become increasingly relevant in the last decades. A relatively recent general methodology, which has been proven very successful, encodes various properties of real-world systems by such formulas, e.g. the correctness of programs [1, 2, 3, 4] or security issues [5]. To check whether the encoded properties hold or not, we need suitable software tools that can check the satisfiability of these encodings in a fully automated manner.

Besides general-purpose computer algebra systems, this functionality is offered by some dedicated *satisfiability modulo theories (SMT) solvers*. These tools use SAT solving to identify *sets of polynomial constraints* whose truth satisfies the Boolean structure of an input formula, and require a *theory solver* that can check the satisfiability of such constraint sets.

There are a few, even though not many, algorithms available that can be implemented to obtain such a theory solver. Incomplete methods like interval constraint propagation [6], virtual substitution [7], and subtropical satisfiability [8] are complemented by the complete method of the *cylindrical algebraic decomposition (CAD)* [9] and a recent adaption of it named the *cylindrical algebraic covering (CALC)* [10]. Employing CALC as a theory solver in traditional SMT solving clearly improves over previous solving approaches regarding computational effort,

SMT'23: 21st International Workshop on Satisfiability Modulo Theories, July 05–06, 2023, Rome, Italy

*Corresponding author.

✉ philipp.baer@rwth-aachen.de (P. Bär); nalbach@cs.rwth-aachen.de (J. Nalbach); abraham@cs.rwth-aachen.de (E. Ábrahám); wcbrown@usna.edu (C. W. Brown)

🆔 0000-0002-2641-1380 (J. Nalbach); 0000-0002-5647-6134 (E. Ábrahám)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

as demonstrated by the cvc5 solver [11], which has won the 2022 SMT competition [12] in the category of QFNRA using the CAIC algorithm.

The QFNRA satisfiability problem is decidable [13] but hard. Both CAD and CAIC might need doubly exponential effort in the worst case [14] and there are currently no cheaper alternatives, even though QFNRA is known to be solvable in singly exponential time [15]. Thus *optimizations* play an important role to improve applicability and practical relevance.

In this paper we propose such an optimization for the CAIC algorithm: we exploit knowledge about the *strictness of input constraints* in order to reduce the computational effort. For constraints in n variables, the CAIC method iteratively computes subsets of \mathbb{R}^n over which the input constraint set is invariantly unsatisfiable, until either it finds a solution or the computed subsets cover \mathbb{R}^n . Such a covering is composed of a finite number of open and non-open subsets, where each non-open one lies in the boundary of an open one. Previous work like e.g. [16] already exploited the well-known fact that if *all* input constraints are strict then only the *open* subsets need to be considered: if there is a solution in a non-open subset N then there is also an ε -ball of solutions around it, therefore also the open subset whose boundary contains N will contain solutions. Skipping non-open subsets does not simply reduce the number of cases, but it avoids the hard cases, which typically require effortful computations with non-rational real-algebraic numbers. In this work, we go further and show that we can neglect some of the non-open subsets even if *not all* (but some) of the input constraints are strict. This paper has three main contributions:

1. We formalize the above-mentioned optimization for the CAIC method, to save effort in the presence of strict input constraints. Since a decomposition is a special type of covering, our results are transferable to the CAD method.
2. We provide a publicly available implementation of the proposed optimization.
3. We evaluate this implementation to evaluate and compare the modification to the original CAIC method.

Outline: We introduce the CAD and the CAIC methods in Section 2 before we present our optimization in Section 3. We discuss the implementation in the SMT-RAT toolbox [17] and provide experimental results to demonstrate effectiveness in Section 4. Finally, we draw conclusions in Section 5.

2. Preliminaries

Quantifier-free non-linear real algebra (QFNRA) Let \mathbb{N} , $\mathbb{N}_{>0}$, \mathbb{Q} and \mathbb{R} denote the set of all natural, natural excluding zero, rational resp. real numbers. For $s \in \mathbb{R}$, its *sign* $sgn(s)$ is 1 if $s > 0$, 0 if $s = 0$, and -1 otherwise. For a set S , we define $\mathcal{P}(S) = \{S' \mid S' \subseteq S\}$.

Assume for the rest of the paper $n \in \mathbb{N}_{>0}$ and some statically ordered *variables* $x_1 < \dots < x_n$. Let $i \in \mathbb{N}_{>0}$ with $i \leq n$. By $\mathbb{Q}[x_1, \dots, x_i]$ we denote the set of all *polynomials* with variables x_1, \dots, x_i and rational coefficients. The *variety* of $p \in \mathbb{Q}[x_1, \dots, x_i]$ is the set of its *real zeros* or *roots* $\{s \in \mathbb{R}^i \mid p(s) = 0\}$. A (QFNRA) *constraint* has the form $p \sim 0$ with $p \in \mathbb{Q}[x_1, \dots, x_n]$ and $\sim \in \{<, \leq, =, \neq, \geq, >\}$; $p \sim 0$ is *strict* iff $\sim \in \{<, >, \neq\}$. A (QFNRA) *formula* φ is a Boolean combination of constraints. By $\varphi(x_1, \dots, x_i)$ we express that $\mathbb{Q}[x_1, \dots, x_i]$ includes

all polynomials appearing in φ . For $(s_1, \dots, s_i) \in \mathbb{R}^i$, by $\varphi(s_1, \dots, s_i, x_{i+1}, \dots, x_n)$ we denote the formula $\varphi[s_1/x_1] \dots [s_i/x_i]$ after substituting s_j for x_j for $j = 1, \dots, i$.

A *cell* is a non-empty, connected subset of \mathbb{R}^i for some $1 \leq i \leq n$. Let $S \subseteq \mathbb{R}^i$ be a cell. Given $\varepsilon \in \mathbb{R}$, $\varepsilon > 0$ and $s \in \mathbb{R}^i$, the ε -ball around s is the set $B_\varepsilon(s) := \{s' \in \mathbb{R}^i \mid |s - s'| \leq \varepsilon\}$. S is *open* iff for all $s \in S$ there is an $\varepsilon \in \mathbb{R}$, $\varepsilon > 0$ such that $B_\varepsilon(s) \subseteq S$. S is *closed* iff $\mathbb{R}^i \setminus S$ is open. The *closure* $cl(S)$ of S is its smallest closed superset, its *interior* $int(S)$ is its largest open subset, and its *boundary* $bound(S)$ is its closure without its interior.

S is *semi-algebraic* if it is the solution set of a QFNRA formula. Assuming $j \in \mathbb{N}_{>0}$ with $1 \leq j < i$, the *projection* of S to \mathbb{R}^j is $S \downarrow_j = \{(s_1, \dots, s_j) \mid (s_1, \dots, s_i) \in S\}$; note that if S is semi-algebraic then $S \downarrow_j$ is semi-algebraic.

S is *sign-invariant* for $P \subseteq \mathbb{Q}[x_1, \dots, x_i]$ if $sgn(p(s)) = sgn(p(s'))$ for all $p \in P$ and all $s, s' \in S$. S is *truth-invariant* (respectively *UNSAT*) for a formula $\varphi(x_1, \dots, x_i)$ if $\varphi(s)$ and $\varphi(s')$ simplify to the same truth value (respectively to **False**) for all $s, s' \in S$. For $P \subseteq \mathbb{Q}[x_1, \dots, x_i]$ and $s \in \mathbb{R}^i$, by $S(P, s)$ we denote the maximal cell $S \subseteq \mathbb{R}^i$ that is sign-invariant for P and contains s .

Cell properties are generalized to sets of cells by requiring the respective property for each cell in the set. E.g., a cell set $D \subseteq \mathcal{P}(\mathbb{R}^i)$ is semi-algebraic if each $S \in D$ is semi-algebraic. We also extend the projection to sets of cells $D \subseteq \mathcal{P}(\mathbb{R}^i)$ as $D \downarrow_j = \{S \downarrow_j \mid S \in D\}$.

Cylindrical algebraic decomposition (CAD) The *CAD method* [9] was introduced by Collins in 1975. Despite its doubly exponential complexity, all complete algorithms in SMT solvers for QFNRA are based on the techniques underlying the CAD. This also holds true for the CALC method. For a formula $\varphi(x_1, \dots, x_n)$, the CAD method decomposes \mathbb{R}^n into a finite set D of cells that are sign-invariant for the polynomials in φ and thus truth-invariant for φ , such that we can decide φ 's satisfiability by evaluating it at one sample point from each $S \in D$.

Definition 1 (Cylindrical algebraic decomposition). Assume $i \in \mathbb{N}_{>0}$ with $i \leq n$.

- A decomposition of \mathbb{R}^i is a finite set $D \subset \mathcal{P}(\mathbb{R}^i)$ such that $\bigcup_{S \in D} S = \mathbb{R}^i$ and either $S = S'$ or $S \cap S' = \emptyset$ for every $S, S' \in D$.
- A decomposition D of \mathbb{R}^i is algebraic if each $S \in D$ is a semi-algebraic cell.
- A decomposition D_i of \mathbb{R}^i is cylindrical if either $i = 1$, or $i > 1$ and $D_{i-1} = \{S \downarrow_{i-1} \mid S \in D\}$ is a cylindrical decomposition of \mathbb{R}^{i-1} .

For a finite set of polynomials $P \subset \mathbb{Q}[x_1, \dots, x_n]$, the CAD method computes CADs D_1, \dots, D_n of $\mathbb{R}^1, \dots, \mathbb{R}^n$ with $D_i \downarrow_{i-1} = D_{i-1}$ for $i = 2, \dots, n$, in two phases:

(1) The *projection phase* computes for $i = n, \dots, 1$, in this order, finite sets of polynomials $P_i \subset \mathbb{Q}[x_1, \dots, x_i]$ such that the varieties of $\bigcup_{j=1}^i P_j$ define the boundaries of the cells in D_i . Starting from $P_n = P$, for $i = n-1, \dots, 1$ we obtain P_i by a *projection operator* $proj_i : \mathcal{P}(\mathbb{Q}[x_1, \dots, x_{i+1}]) \rightarrow \mathcal{P}(\mathbb{Q}[x_1, \dots, x_i])$ that eliminates the highest variable from its input polynomials. We do not define $proj_i$ here, it suffices to assume that it maintains some properties such that the resulting CAD is sign-invariant for P .

(2) In the *lifting phase*, compute the cells in the CADs D_1, \dots, D_n , in this order. Instead of representing them explicitly, usually a sample point for every cell is generated. For $i \in \{1, \dots, n\}$ and $s \in \mathbb{R}^{i-1}$ (with $\mathbb{R}^0 = \{\emptyset\}$), let $\Xi_s = \{\xi_{s,1}, \dots, \xi_{s,k_s}\}$ be the ordered set of all real roots of the

univariate polynomials from $\{p(s, x_i) \mid p \in P_i\}$. Let furthermore $I_s = \{(-\infty, \infty)\}$ if $\Xi_s = \emptyset$, and otherwise let I_s consist of all point-sets $\{\xi_{s,j}\}$, $j = 1, \dots, k_s$ (which cover the *sections*, i.e. the cells consisting of roots of P_i), and the intervals $(-\infty, \xi_1), (\xi_1, \xi_2), \dots, (\xi_{k_s-1}, \xi_{k_s}), (\xi_{k_s}, \infty)$ (which cover the *sectors*, i.e. the open cells between the sections). Assume furthermore a fixed but arbitrary function *sample* that assigns to each non-empty interval a value from this interval. For $i = 1, \dots, n$, based on S_{i-1} with $S_0 = \{()\}$, we iteratively define the sample sets for D_i as $S_i = \{(s, s_i) \mid s \in S_{i-1} \wedge I \in I_s \wedge s_i = \text{sample}(I_s)\}$.

Example 1. Consider $P = P_2 = \{p_1 : -x_1^2 - x_2 + 1, p_2 : x_1^2 - x_2 - 1, p_3 : (x_1 - 0.5)^2 + (x_2 + 1.5)^2 - 0.25, p_4 : x_1 + 0.5\}$ with projection $P_1 = \text{proj}_2(P_2)$. Figure 1 depicts the varieties of the polynomials from P_2 , and below it those from P_1 . The intervals defined by the real zeros from P_1 form a sign-invariant CAD D_1 for P_1 . The sign-invariant CAD D_2 for P_2 is cylindrical over D_1 , i.e. the cells of D_2 are arranged in cylinders over D_1 . In D_1 , the sections are the cells containing only a real root from P_1 , the open intervals are the sectors. In D_2 , we extend the varieties with the cylinder boundaries induced by D_1 , and define D_2 's sections as the intersection points of lines as well as the line segments bounded by them, and the sectors as the open cells bounded by the sections. \square

We note that not only the computation of the projection is computationally expensive, but also the number of samples grows exponentially with the number of variables. Although lifting over rational samples (that involves plugging in the sample point to obtain univariate polynomials and isolating their real roots) is rather efficient, in general, expensive lifting over non-rational *real-algebraic* samples (which requires similar machinery as the projection) cannot be avoided.

Cylindrical algebraic covering CADs are finer than what we actually need for checking the satisfiability of QFNRA formulas, as we either need to find a satisfying sample, or *cover* \mathbb{R}^n with UNSAT cells. Example 1 illustrates that cells which are UNSAT for a constraint are often split into several smaller cells due to sign changes of some other polynomials. To avoid such splits, the *cylindrical algebraic covering* (CAIC) [10] method, which uses the techniques from the CAD, relaxes sign-invariance for truth-invariance, still keeping the cylindrical arrangement of UNSAT cells but allowing overlaps between them. Here we give a brief overview on the CAIC method for *conjunctions* of constraints, and refer to [10, 18] for the general case.

Definition 2 (Cylindrical algebraic covering). Assume $i \in \mathbb{N}_{>0}$ with $i \leq n$.

- A covering of \mathbb{R}^i is a finite set $C \subseteq \mathcal{P}(\mathbb{R}^i)$ such that $\bigcup_{S \in C} S = \mathbb{R}^i$.
- A covering C of \mathbb{R}^i is algebraic if each $S \in C$ is a semi-algebraic cell.
- A covering C of \mathbb{R}^i is cylindrical if either $i = 1$, or $i > 1$ and $C_{i-1} = \{S \downarrow_{i-1} \mid S \in C_i\}$ is a cylindrical covering of \mathbb{R}^{i-1} .

The CAIC method is *sample-guided*: in contrast to CAD, it does not start with projection but with a dimension-wise guess of values for a sample, with the aim to satisfy the input formula φ . For a current partial sample $(s_1, \dots, s_{i-1}) \in \mathbb{R}^{i-1}$ with $1 \leq i \leq n$, we iteratively identify intervals $I \subseteq \mathbb{R}$ such that for any $s_i \in I$, $\varphi(s_1, \dots, s_i)$ is unsatisfiable. We continue this process until either we find a solution (along with a partial UNSAT covering) or the intervals




Figure 1: Examples for CAD (left) and CAIC (right). *Left:* Varieties of P_2 (top) and P_1 (bottom) from Example 1. Cells of the sign-invariant CADs are coloured blue. *Right:* Varieties of P_2 (top) and P_1 (bottom) from Example 2. Inwards-pointing arrows mark unsatisfiable areas for φ 's constraints (distinguished by different colours). UNSAT-intervals generated during the CAIC computations are shown as bracket-bounded thick black lines (open intervals) and short line segments (point intervals). UNSAT-cells for φ are coloured blue. The one-dimensional samples $s = 0.25$ and $s = 1$ are shown as thin vertical lines.

cover the whole \mathbb{R} . In the latter case, if $i = 1$ then the problem is unsatisfiable and a complete CAIC is returned; otherwise, we backtrack one level to sampling for x_{i-1} , generalize s_{i-1} to an unsatisfying interval, and try to guess another value for x_{i-1} outside the already excluded intervals. The interval generalization I of s_{i-1} is computed using real root isolation and a reduced version of the CAD projection. Intuitively, it contains s_{i-1} and other points s'_{i-1} for which $\varphi(s_1, \dots, s_{i-2}, s'_{i-1})$ is unsatisfiable for same reason as $\varphi(s_1, \dots, s_{i-2}, s_{i-1})$.

In general, CAICs require less computational effort than CADs: the sample-based search in CAIC saves parts of the projection executed in CAD, as well as the lifting effort corresponding to them.

Example 2. Re-using $P = \{p_1 : -x_1^2 - x_2 + 1, p_2 : x_1^2 - x_2 - 1, p_3 : (x_1 - 0.5)^2 + (x_2 + 1.5)^2 - 0.25, p_4 : x_1 + 0.5\}$ from Example 1, we consider the formula $\varphi := p_1 < 0 \wedge p_2 > 0 \wedge p_3 \geq 0 \wedge p_4 \geq 0$. Figure 1 illustrates on the right the following CAIC computations.

Samples $x_1 \in (-\infty, -0.5)$ violate $p_4 \geq 0$. Outside $(-\infty, -0.5)$ we pick $s_1 = 0.25$ for x_1 , and consider $\varphi(0.25, x_2)$. The constraint $p_1(0.25, x_2) < 0$ is unsatisfiable for $x_2 \in (-\infty, \frac{15}{16})$, and $p_2(0.25, x_2) > 0$ is unsatisfiable for $x_2 \in (-\frac{15}{16}, \infty)$, together covering the real line. We generalize the UNSAT result to the cell $(-1, 1)$ containing $s_1 = 0.25$ in the D_1 CAD for $\{p_1, p_2\}$ (see also Example 3).

Outside $(-\infty, -0.5) \cup (-1, 1)$ we pick $s_1 = 1$ for x_1 , and compute a covering of unsatisfying intervals for $\varphi(1, x_2)$ as depicted in the two-dimensional coordinate system. This time the sample $s_1 = 1$ is a real root, which cannot be generalized further than the section $[1, 1]$ for x_1 .

Next, we pick a satisfying sample with $s_1 = 1.5$ and $s_2 = 0$, and the algorithm terminates.

This example shows where the CAIC method is more efficient than the CAD: Firstly, single constraints like $p_4 \geq 0$ above can be used to rule out parts of the search space requiring fewer projection and lifting steps. Secondly, the CAIC method does not necessarily involve all polynomials in projection and corresponding lifting steps if they are redundant, as e.g. $p_3 \geq 0$ above excludes part of the search space that is already ruled out by the other constraints. \square

Each UNSAT interval $I \subseteq \mathbb{R}$ generated during the CAIC computations over some partial sample $s \in \mathbb{R}^{i-1}$ corresponds to an UNSAT cell $S \subseteq \mathbb{R}^i$ with $\{s\} \times I = (\{s\} \times \mathbb{R}) \cap S$. This UNSAT cell is represented *implicitly* by a set of polynomials $P \subset \mathbb{Q}[x_1, \dots, x_i]$ and the sample (s, s_i) for some $s_i \in I$ such that, except for some special cases, S is the maximal sign-invariant cell for P that contains (s, s_i) , i.e. $S = S(P, (s, s_i))$.

For generalizing a covering of intervals $I_1, \dots, I_k \subseteq \mathbb{R}$ (i.e. $\cup_{j=1}^k I_j = \mathbb{R}$) over some sample point $s \in \mathbb{R}^{i-1}$, the CAIC method defines a partial projection operator that projects the implicitly represented cells to an $(i-1)$ -dimensional cell containing s as follows:

Definition 3. The covering projection operator proj_{cov} is a function which, for any $i, k \in \mathbb{N}_{>0}$, $P_1, \dots, P_k \subset \mathbb{Q}[x_1, \dots, x_i]$, $s \in \mathbb{R}^{i-1}$ and $s'_1, \dots, s'_k \in \mathbb{R}$ with $s \times \mathbb{R} \subseteq \cup_{j=1}^k S(P_j, (s, s'_j))$ as input, returns a set of polynomials $P = \text{proj}_{\text{cov}}(P_1, \dots, P_k, s, s'_1, \dots, s'_k) \subseteq \mathbb{Q}[x_1, \dots, x_{i-1}]$ with the property that $S(P, s) \times \mathbb{R} \subseteq \cup_{j=1}^k S(P_j, (s, s'_j))$.

The correctness of an UNSAT covering is given by the following theorem:

Theorem 1. Let φ be a formula in variables x_1, \dots, x_n , $i > 1$, $P_1, \dots, P_k \subseteq \mathbb{Q}[x_1, \dots, x_i]$, $s \in \mathbb{R}^{i-1}$ and $s'_1, \dots, s'_k \in \mathbb{R}$ such that $s \times \mathbb{R} \subseteq \cup_{j=1}^k S(P_j, (s, s'_j))$ and for all $j = 1, \dots, k$ the cell $S(P_j, (s, s'_j))$ is UNSAT for φ .

Then $S(\text{proj}_{\text{cov}}(P_1, \dots, P_k, s, s'_1, \dots, s'_k), s)$ is UNSAT for φ .

3. Exploiting the Strictness of Constraints

Strict constraints $p < 0$ are never satisfied at the real roots of p . Therefore, when checking the satisfiability of a formula φ that contains only strict constraints (and no negations), sections can be neglected in CAD and CAIC computations, i.e. real roots can be omitted during sample construction. This observation is exploited e.g. in [16].

In this work we propose an approach that allows to omit real roots during sample construction in certain cases even if the input formula contains also weak constraints. We start with illustrating the idea on an example.

Example 3. In Example 2, for each $s_1 \in \mathbb{R}$, both $p_1(s_1, x_2)$ and $p_2(s_1, x_2)$ have exactly one real root, which we denote by $\xi_1(s_1)$ respectively $\xi_2(s_1)$.

For the sample $s_1 = 0.25$ for x_1 , we covered x_2 by the intervals $(-\infty, \frac{15}{16})$ violating $p_1 < 0$, and $(-\frac{15}{16}, \infty)$ violating $p_2 > 0$. The above intervals represent the cells $\{(s_1, s_2) \in \mathbb{R}^2 \mid s_2 < \xi_1(s_1)\}$ and $\{(s_1, s_2) \in \mathbb{R}^2 \mid s_2 > \xi_2(s_1)\}$. The generalization of $s_1 = 0.25$ is the maximal interval $(-1, 1)$ over which the above cells still build a covering (see Figure 1).

Note that $s_2 = \xi_1(s_1)$ implies $p_1 = 0$, and $s_2 = \xi_2(s_1)$ implies $p_2 = 0$. Thus, the closures $\{(s_1, s_2) \in \mathbb{R}^2 \mid s_2 \leq \xi_1(s_1)\}$ and $\{(s_1, s_2) \in \mathbb{R}^2 \mid s_2 \geq \xi_2(s_1)\}$ are still unsatisfying for $p_1 < 0$

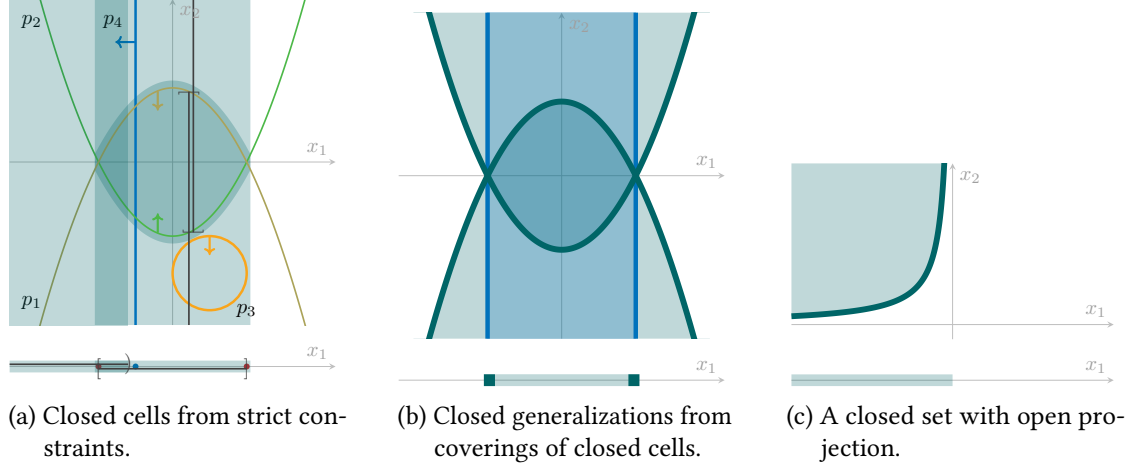


Figure 2: Projections of closed cells.

resp. $p_2 > 0$. Now, since both cells are closed, they cover x_2 over the closed generalization $[-1, 1]$, the covering intervals at $s_1 = \pm 1$ being $(-\infty, 0]$ and $[0, \infty)$, which makes it unnecessary to consider the sample $s_1 = 1$ from Example 2. \square

Clearly, if we are able to deduce closed intervals, the computed CALC consists not only of fewer cells, but we can also avoid computationally intensive lifting operations over potentially non-rational algebraic numbers. The following example demonstrates that for two-dimensional formulas we could go even further.

Example 4. Assume in the previous example that the first constraint would be non-strict, then for the constraints $p_1(0.25, x_2) \leq 0$ and $p_2(0.25, x_2) > 0$ we would achieve the covering $(-\infty, \frac{15}{16})$ and $[-\frac{15}{16}, \infty)$ for x_2 . Now, only the second cell is closed. However, this still suffices to cover $[-1, 1]$ also at its endpoints $s_1 = \pm 1$ by $(-\infty, 0)$ and $[0, \infty)$. \square

However, generalizing this observation to constraint sets with more than two variables is non-trivial; we therefore focus on the case where *all* intervals of a covering are closed.

To that end, we change our view from UNSAT intervals to the UNSAT cells they represent. During the CALC computations, some UNSAT cells violate a constraint (as indicated by the arrows in Figure 2a); as illustrated in Example 3, we can *close* those cells S that violate *strict* constraints $p \sim 0$ (i.e. $p(s) \not\sim 0$ for all $s \in S$) without losing the UNSAT property of the cell (i.e. $p(s) \not\sim 0$ for all $s \in cl(S)$).

Furthermore, for any covering by UNSAT intervals which represent the i -dimensional *closed* cells $S_1, \dots, S_k \subseteq \mathbb{R}^i$, if $S \subseteq \mathbb{R}^{i-1}$ is a possible generalization of the current sample (i.e. $S \times \mathbb{R} \subseteq \cup_{j=1}^k S_j$), then also the closure of S is a valid generalization ($cl(S) \times \mathbb{R} \subseteq \cup_{j=1}^k S_j$). This fact is formalized in the following theorem and visualized in Figure 2b.

Theorem 2. Assume $i, k \in \mathbb{N}_{>0}$, closed cells $S_1, \dots, S_k \subseteq \mathbb{R}^i$, and a cell $S \subseteq \mathbb{R}^{i-1}$ such that $S \times \mathbb{R} \subseteq \cup_{j=1}^k S_j$. Then $cl(S) \times \mathbb{R} \subseteq \cup_{j=1}^k S_j$.

Proof. Assume for contradiction that there exists an $s \in bound(S) \times \mathbb{R} \subseteq \mathbb{R}^i$ such that $s \notin \cup_{j=1}^k S_j$. Then $s \notin S_j$, i.e. $s \in \mathbb{R}^i \setminus S_j$ for all $j \in \{1, \dots, k\}$. As the sets S_j are closed, their

complements $\mathbb{R}^i \setminus S_j$ are open. By definition of an open cell, for each $j \in \{1, \dots, k\}$ there exist $\varepsilon_j > 0$ such that $B_{\varepsilon_j}(s) \subseteq \mathbb{R}^i \setminus S_j$. Let ε be the smallest such ε_j under all $j \in \{1, \dots, k\}$. Then $B_\varepsilon(s) \subseteq \mathbb{R}^i \setminus S_j$ for all $j \in \{1, \dots, k\}$. As $s \in \text{bound}(S) \times \mathbb{R} = \text{bound}(S \times \mathbb{R})$, by the definition of $B_\varepsilon(s)$ there exist an $s' \in B_\varepsilon(s) \cap (S \times \mathbb{R})$, in other words, $s' \notin S_j$ for all $j \in 1, \dots, k$, which is a contradiction to the assumption that $S \times \mathbb{R} \subseteq \bigcup_{j=1}^k S_j$. \square

The above theorem might seem straight-forward at the first sight, but there are some special cases that make it less trivial, e.g. that the projection of a closed cell might not be closed as shown in Figure 2c.

Now, we apply the general Theorem 2 from above to Theorem 1 to obtain a variant that supports the derivation of closed cells.

Theorem 3. *Let φ be a formula in variables x_1, \dots, x_n , $i > 1$, $P_1, \dots, P_k \subseteq \mathbb{Q}[x_1, \dots, x_i]$, $s \in \mathbb{R}^{i-1}$ and $s'_1, \dots, s'_k \in \mathbb{R}$ such that $s \times \mathbb{R} \subseteq \bigcup_{j=1}^k \text{cl}(S(P_j, (s, s'_j)))$ and for all $j = 1, \dots, k$ the cell $\text{cl}(S(P_j, (s, s'_j)))$ is UNSAT for φ .*

Then $\text{cl}(S(\text{proj}_{\text{cov}}(P_1, \dots, P_k, s, s'_1, \dots, s'_k), s))$ is UNSAT for φ .

Proof sketch. Let $l \geq k$, $P_{k+1}, \dots, P_l \subseteq \mathbb{Q}[x_1, \dots, x_i]$, $s'_{k+1}, \dots, s'_l \in \mathbb{R}$ such that $s \times \mathbb{R} \subseteq \bigcup_{j=1}^l S(P_j, (s, s'_j))$ and for every $j' = k+1, \dots, l$ there exists a $j \in \{1, \dots, k\}$ such that $P_{j'} = P_j$ and $(s, s'_{j'}) \in \text{bound}(S(P_j, (s, s'_j)))$, that means $S(P_{j'}, (s, s'_{j'})) \subseteq \text{bound}(S(P_j, (s, s'_j)))$.

Then by Theorem 1 it holds for $P := \text{proj}_{\text{cov}}(P_1, \dots, P_l, s, s'_1, \dots, s'_l)$ that $S(P, s) \times \mathbb{R} \subseteq \bigcup_{j=1}^l S(P_j, (s, s'_j)) = \bigcup_{j=1}^k \text{cl}(S(P_j, (s, s'_j)))$.

We now apply Theorem 2 to $S(P, s)$ and $\text{cl}(S(P_j, (s, s'_j)))$, $j = 1, \dots, k$ and obtain $\text{cl}(S(P, s) \times \mathbb{R}) \subseteq \bigcup_{j=1}^k \text{cl}(S(P_j, (s, s'_j)))$.

For the theorem, it remains to show that $P = \text{proj}_{\text{cov}}(P_1, \dots, P_k, s, s'_1, \dots, s'_k)$. A formal proof would require the definition of the details of the projection operator proj_{cov} , which we had to omit due to space restrictions. At this point, we just state without proving that adding the additional cells does not change the projection. We justify that as each additional cell $S(P_{j'}, s'_{j'})$, $j = k+1, \dots, l$ describes the boundary of a neighbouring open cell $S(P_j, s'_j)$ for some $j \in \{1, \dots, k\}$ and the defining polynomial sets $P_j = P_{j'}$ are equal, and thus the ‘skeleton’ of the covering does not change. \square

This result yields a simple adaption of the CAIC algorithm: Along with each implicit cell representation, we store a Boolean flag that indicates whether the corresponding unsatisfiable cell is closed or not. If the flag is set, the corresponding cell is closed, thus we can set the bounds of the witnessing interval to closed. Whenever we compute the base cell of a covering consisting of closed cells, we can easily deduce the flag for the new cell (it is true whenever the parents’ flags are all true). Thus, any implementation of the CAIC implementation can be adapted for the theorem by only superficial changes in the code.

We provide a 3D example for the covering method and our adaption in the appendix of [19].

4. Experimental results

We implemented the proposed method to exploit strict constraints in our SMT-RAT [17] solver, using standard preprocessing and DPLL(T) solving with an implementation of the CAIC method

Table 1

Number of solved instances on the whole benchmark set.

Solver	SAT	UNSAT
CA1C	4553	4625
CA1C-I	4610	4648
CA1C-IH	4609	4648
Total	5069	5379
	(1104 unknown)	




Figure 3: Number of instances solved by CA1C-I and their (maximal) ratio of flagged/closed cells.


as the only theory solver. The implementation is accessible at <https://doi.org/10.5281/zenodo.7900518>. We execute the solver on the *QF_NRA* benchmark library from SMT-LIB [20] (as of April 2022), consisting of 11552 instances that stem from 11 different families. Each formula is solved on a CPU with 2.1 GHz with a timeout of 60 seconds and a memory limit of 4 GB. In the following, we denote the original CA1C solver by CA1C and the modification that maintains interval flags to indicate closed cells by CA1C-I.

Table 1 shows the overall performance of CA1C and CA1C-I. The modified method CA1C-I solves 80 instances more than CA1C, whereby more than two thirds of this gain is on SAT instances, the remaining on UNSAT instances. CA1C times out on 77 instances which are solved by CA1C-I. Conversely, CA1C-I times out on 7 instances which are solved by CA1C.

Applications of the theorem Figure 3 depicts the share of the *derived* implicit cell representations that carry a True flag (i.e. Theorem 3 was applicable during its creation). More than half of the instances are already solved by the SAT solver (the theory solver is never called) or the call to the CA1C contains only univariate polynomials. The theorem is never applied on 2518 instances. The theorem is applied at least once in 1699 instances. The theorem is always applied on 1162 instances (i.e. all cells are closed).

We focus the further analysis on the interesting instances where the theorem was applied at least once: the instances solved by both solvers and the theorem was applied at least once (1623 instances), instances solved only by CA1C (8 instances), and instances solved only by CA1C-I (88 instances).

Running times Figure 4a compares the running times of CA1C and CA1C-I on the interesting instances. The bottom left cluster of instances is computationally easy; small deviations in running time between CA1C and CA1C-I are negligible. Most instances with higher running time are located near the equality line, i.e. CA1C and CA1C-I are equally efficient on them. There are 30 instances solved faster by more than 0.1 seconds on CA1C than on CA1C-I (including instances solved only by CA1C, see area below the equality line). The other way around, 197



(a) Running times of CA1C and CA1C-I.

(b) Samples of CA1C and CA1C-I

Figure 4: Scatter plots for running time and number of samples. \perp denotes a timeout on an instance.

instances are solved faster by CA1C-I (see area above the equality line). We conclude that CA1C-I can significantly improve the running time on certain instances.

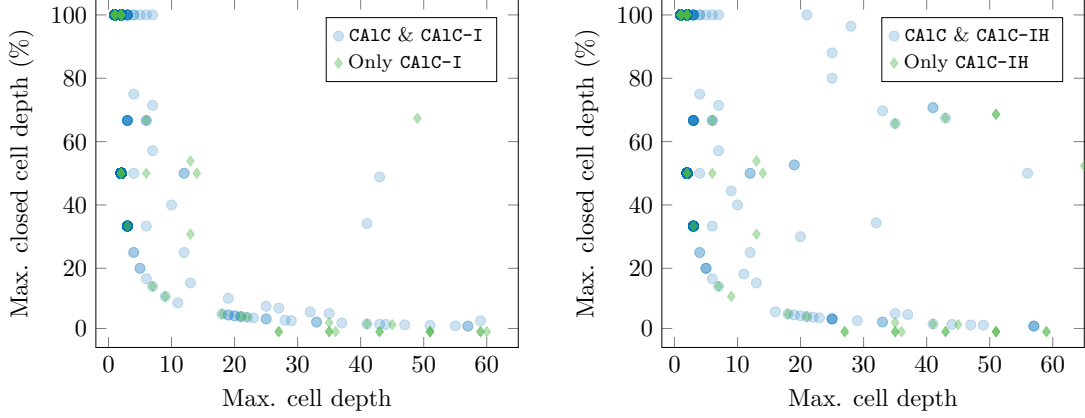
In particular, some UNSAT instances (blue crosses) deviate from the equality line for running times greater than one second. CA1C-I is able to cover the entire space with fewer UNSAT cells than CA1C; this effect is less significant on the SAT instances, where only a partial covering is computed.

Number of samples We expect that the advantage of the proposed optimization is due to the fact that the application of Theorem 3 reduces the number of samples constructed in the CA1C method.

To evaluate this correlation, Figure 4b compares the number of partial sample points of CA1C and CA1C-I. Though the number of samples is often similar, CA1C-I tends to generate fewer samples especially on the larger instances.

Iterative applications of the theorem As Theorem 3 can only be applied if *all* cells forming the covering are closed, the question raises how often it can be applied iteratively. To that end, we say that the cells forming a covering are the *parents* of the covering's base cell. In that sense, we define for every cell its *depth* as the distance to its 'oldest' ancestor. Figure 5a plots for every instance the maximal depth among all cells versus the relative maximal depth among all closed cells.

We clearly see that the theorem is mostly applicable to cells with low depth. In particular, the 1308 instances at the top left of the plot, where at least one closed cell has maximal depth, have a total maximal depth below 10. Further, the instances where the theorem is only applied to cells with depth one form a visible hyperbola, which are 1159 of the 1711 depicted instances. We conclude that the performance gains stem mostly from 'superficial' application of the theorem. Still, there are some instances above this hyperbola, representing non-trivial applications of our theorem.



(a) Cell depth ratio of CA1C-I

(b) Cell depth ratio of CA1C-IH

Figure 5: Maximal cell depth and the relative maximal closed cell depth.

Modification of the covering heuristic The modification CA1C-I is clearly more efficient than CA1C; as we just observed, these gains are due to ‘superficial’ application of Theorem 3. We now aim to adapt the CA1C-I method to support the application of the theorem also to cells of higher depth. Whenever the CA1C method finds a covering of cells at some sample, the choice of the cells forming the covering is not unique; redundancies in the formula might allow for a choice of the cells. CA1C-I heuristically minimizes the number of cells forming a covering. We propose an adaption CA1C-IH which first tries to cover using closed cells only, and falls back to the default heuristic if it fails.

Figure 5b shows that the theorem is now applied on more instances with higher depth. However, the running times do not improve significantly (not shown here); Table 1 shows that even one instance less is solved. Thus, more sophisticated heuristics are desirable which find a better trade off between a covering of ‘good’ intervals and supporting the theorem applicability.

5. Conclusion

The cylindrical algebraic covering method admits reducing the number of projection and lifting operations compared to the cylindrical algebraic decomposition switching from being sign-invariant for a set of polynomials to being truth-invariant for an input formula. In this paper, we propose a natural extension to the CA1C method that exploits strict constraints in the input formula: If a strict constraint is unsatisfiable in some cell, then it is so in the closure of that cell. Our adaption allows to carry this information through the CA1C algorithm to avoid lifting over roots of polynomials, which is desirable as this is usually computationally expensive.

The proposed adaption is easy and efficient to implement. Our experimental evaluation concluded that (1) we gain a good portion of newly solved instances, (2) these gains are due to reduced number of lifting steps, and (3) our modification is still ‘superficial’ and leaves potential for future investigation into the topic.

Future work consists of advanced theoretical work as motivated by Example 4, and better heuristics for choosing good coverings as motivated in the last paragraph of Section 4.

Acknowledgments

Jasper Nalbach was supported by the DFG RTG 2236/2 *UnRAVeL*. We thank James Davenport and Matthew England for fruitful discussions.

References

- [1] P. Berger, J.-P. Katoen, E. Ábrahám, M. T. B. Waez, T. Rambow, Verifying auto-generated C code from Simulink, in: Proc. of FM'18, Springer, 2018, pp. 312–328. doi:10.1007/978-3-319-95582-7_18.
- [2] D. Kroening, M. Tautschnig, CBMC - C bounded model checker, in: Proc. of the 20th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'14), volume 8413 of LNCS, Springer, 2014, pp. 389–391. URL: https://doi.org/10.1007/978-3-642-54862-8_26. doi:10.1007/978-3-642-54862-8_26.
- [3] M. Heizmann, J. Hoenicke, A. Podelski, Software model checking for people who love automata, in: Proc. of the 25th Int. Conf. on Computer Aided Verification (CAV'13), volume 8044 of LNCS, Springer, 2013, pp. 36–52. URL: https://doi.org/10.1007/978-3-642-39799-8_2. doi:10.1007/978-3-642-39799-8_2.
- [4] A. Cimatti, SMT-based software model checking - Explicit scheduler, symbolic threads, in: D. V. Hung, M. Ogawa (Eds.), Proc. of the 11th Int. Symp. on Automated Technology for Verification and Analysis (ATVA'13), volume 8172 of LNCS, Springer, 2013, p. 23. URL: https://doi.org/10.1007/978-3-319-02444-8_3. doi:10.1007/978-3-319-02444-8_3.
- [5] G. Snelting, Quantifier elimination and information flow control for software security, in: Proceedings of the Algorithmic Algebra and Logic, A3L'05, BoD, 2005, pp. 237–242.
- [6] M. Fränzle, C. Herde, T. Teige, S. Ratschan, T. Schubert, Efficient solving of large non-linear arithmetic constraint systems with complex boolean structure, Journal on Satisfiability, Boolean Modeling and Computation 1 (2006) 209–236. doi:10.3233/SAT190012.
- [7] F. Corzilius, E. Ábrahám, Virtual substitution for SMT-solving, in: Proceedings of the Fundamentals of Computation Theory, FCT'11, Springer, 2011, pp. 360–371. doi:10.1007/978-3-642-22953-4_31.
- [8] P. Fontaine, M. Ogawa, T. Sturm, X. T. Vu, Subtropical satisfiability, in: Proceedings of the Frontiers of Combining Systems, FroCoS'17, Springer, 2017, pp. 189–206. doi:10.1007/978-3-319-66167-4_11.
- [9] G. E. Collins, Quantifier elimination for real closed fields by cylindrical algebraic decomposition, in: Proceedings of the Second GI Conference on Automata Theory and Formal Languages, ATFL'75, Springer, 1975, pp. 134–183. doi:10.1007/978-3-7091-9459-1_4.
- [10] E. Ábrahám, J. H. Davenport, M. England, G. Kremer, Deciding the consistency of non-linear real arithmetic constraints with a conflict driven search using cylindrical algebraic coverings, Journal of Logical and Algebraic Methods in Programming 119 (2021) 100633. doi:10.1016/j.jlamp.2020.100633.
- [11] H. Barbosa, C. W. Barrett, M. Brain, G. Kremer, H. Lachnitt, M. Mann, A. Mohamed, M. Mohamed, A. Niemetz, A. Nötzli, A. Ozdemir, M. Preiner, A. Reynolds, Y. Sheng, C. Tinelli, Y. Zohar, cvc5: A versatile and industrial-strength SMT solver, in: Proc.

- of the 28th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'22), volume 13243 of *LNCS*, Springer, 2022, pp. 415–442. URL: https://doi.org/10.1007/978-3-030-99524-9_24. doi:10.1007/978-3-030-99524-9_24.
- [12] SMT-COMP 2022, <https://smt-comp.github.io/2022/>, 2022.
 - [13] A. Tarski, A Decision Method for Elementary Algebra and Geometry, RAND Corporation, 1951. doi:10.1007/978-3-7091-9459-1_3.
 - [14] J. H. Davenport, J. Heintz, Real quantifier elimination is doubly exponential, *J. Symbolic Comp.* 5 (1988) 29–35.
 - [15] C. W. Brown, J. H. Davenport, The complexity of quantifier elimination and cylindrical algebraic decomposition, in: *Proceedings of ISSAC'07*, 2007, pp. 54–60.
 - [16] C. W. Brown, Open non-uniform cylindrical algebraic decompositions, in: *Proceedings of the 2015 ACM on International Symposium on Symbolic and Algebraic Computation*, ISSAC'15, Association for Computing Machinery, 2015, p. 85–92. doi:10.1145/2755996.2756654.
 - [17] F. Corzilius, G. Kremer, S. Junges, S. Schupp, E. Ábrahám, SMT-RAT: An open source C++ toolbox for strategic and parallel SMT solving, in: *Proceedings of the Theory and Applications of Satisfiability Testing*, SAT'15, Springer, 2015, pp. 360–368. doi:10.1007/978-3-319-24318-4_26.
 - [18] G. Kremer, J. Nalbach, Cylindrical algebraic coverings for quantifiers, in: *Proceedings of the 7th International Workshop on Satisfiability Checking and Symbolic Computation (SC² 2022)*, 2022. URL: <https://nafur.github.io/static/2022-scsc-qe-coverings.pdf>.
 - [19] P. Bär, J. Nalbach, E. Ábrahám, C. W. Brown, Exploiting strict constraints in the cylindrical algebraic covering, 2023. doi:<https://doi.org/10.48550/arXiv.2306.16757>.
 - [20] C. Barrett, P. Fontaine, C. Tinelli, The satisfiability modulo theories library (SMT-LIB), www.SMT-LIB.org, 2021.