Merging Adjacent Cells during Single Cell Construction ⋆

Jasper Nalbach[1[0000−0002−2641−1380]] and Erika Ábrahám[1[0000−0002−5647−6134]]

RWTH Aachen University, Aachen, Germany
nalbach@cs.rwth-aachen.de

Abstract. The *cylindrical algebraic decomposition (CAD)* is currently the only complete method used in practise for answering questions about *real algebra*, despite its doubly exponential complexity. Recently, some novel algorithms like *NLSAT*, *CAlC* and *NuCAD* for *satisfiability checking* respectively *quantifier elimination* have been proposed, which build on the CAD idea to generalize a sample point to a connected set (*cell*) of points that share certain properties with the sample. This process is called *single cell construction*. In this paper, we adapt this method to potentially generate bigger cells by detecting that certain adjacent cells maintain those relevant invariance properties. For formalizing the resulting algorithm in this paper, we generalize the notion of delineability to *local delineability*. An experimental evaluation of a first implementation in *NLSAT* is provided.

1 Introduction

The *cylindrical algebraic decomposition (CAD)* method [5] answers certain questions about *real algebraic* formulas. Despite its doubly exponential complexity, it is currently still the only complete method used in practise. Applications include *quantifier elimination* - as implemented in `QEPCAD`, `Redlog` and a range of commercial tools - and *satisfiability-modulo-theories (SMT) solving* for the theory of *non-linear arithmetic* - as implemented in `z3`, `cvc5`, `yices2`, and `SMT-RAT`.

The CAD method decomposes the real space into finitely many connected sets (called *cells*) such that certain properties of the input polynomials are invariant in each of these sets. Computing such decompositions is expensive and is often finer than needed for answering the desired question. Modern algorithms such as *NLSAT* [6], *CAlC* [1, 7] and *NuCAD* compute only a sequence of coarser decompositions, heuristically reducing the computational effort. These algorithms essentially use a *single cell construction* algorithm (first introduced in [6], extended in [4, 10]), which is based on the CAD idea and computes a single cell

⋆ Jasper Nalbach was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of RTG 2236 *UnRAVeL* and AB 461/9-1 *SMT-ART*. Computations were performed with computing resources granted by RWTH Aachen University under project rwth1560.

that maintains some properties of some polynomials and contains a given sample point.

These algorithms are more efficient than using a full CAD for two reasons: Firstly, during single cell construction, we can identify which parts of the CAD projection polynomials are relevant to the current cell, thus reduce the computational effort by avoiding expensive computations. Secondly, these algorithms exploit the Boolean structure from the input formula and sign conditions on the polynomials to determine which parts of the search space need to be excluded by the single cell construction. Larger cells facilitate these mechanisms, as they exclude larger parts of the search space. Further, for some applications like quantifier elimination, not only the efficient computation is relevant, but also the quality of the computed results; larger cells may reflect the problem's structure better and thus be more useful to the user.

Existing single cell construction algorithms already compute cells that are the merger of adjacent cells computed by a CAD by reducing the amount of projection polynomials. This paper carries these ideas forward, and merges adjacent cells based on the shape of the varieties. Since the early years of CAD, there is the idea of determining *adjacencies* of cells (also from different cylinders) and *clustering*, where the number of samples during lifting is reduced. We clarify that in this paper, we only consider adjacencies in the same cylinder. We thus use the term *merging* to distinguish our approach from clustering.

Contribution. We extend the levelwise single cell construction presented in [10], which we introduce in the preliminaries in Section 2. In Section 3, we present our idea and how to detect computationally the cells that can be merged. In Section 4, we formalize these ideas by introducing a new notion of *local delineability* and give the necessary theorems for computing the projection. Afterwards, we further refine these ideas in Section 5. Finally, we provide an experimental evaluation of our first implementation in the context of *NLSAT* in Section 6 and conclude in Section 7.

2 Preliminaries

We briefly introduce our notation and refer to the preliminaries of [10] for details.

Let $\mathbb{N}$, $\mathbb{N}_{>0}$, $\mathbb{Q}$, and $\mathbb{R}$ denote the set of all natural (incl. 0), positive integer, rational, and reals numbers respectively. For $i, j \in \mathbb{N}$ with $i < j$, we define $[i..j] = \{i, \dots, j\}$ and $[i] = [0..i]$. For $i, j \in \mathbb{N}_{>0}, j \leq i$ and $r \in \mathbb{R}^i$, we denote by r_j the j-th component of r and by $r_{[j]}$ the vector $(r_1, \dots, r_j)$. Let $f, g : D \to E$ and let $<$ be a total order on E. We write $f < g$ if $f(d) < g(d)$ for all $d \in D$ and $f \leq g$ if $f(d) \leq g(d)$ for all $d \in D$.

We work with the *variables* $x_1, \dots, x_n$ with $n \in \mathbb{N}_{>0}$ under a fixed *ordering* $x_1 \prec x_2 \prec \dots \prec x_n$. A *polynomial* is built from a set of variables and numbers from $\mathbb{Q}$ using addition and multiplication. We use $\mathbb{Q}[x_1, \dots, x_i]$ to denote *multivariate* polynomials in those variables. A polynomial p *is of level* j ($\text{level}(p) = j$) if x_j is the largest variable in p with non-zero coefficient.

Let $i \in [n]$ and $p, q \in \mathbb{Q}[x_1, \dots, x_i]$ of level i. For $j \in [1..i]$ and $r = (r_1, \dots, r_j) \in \mathbb{R}^j$ we write $p(r, x_{j+1}, \dots, x_i)$ for the polynomial p after substituting $r_1, \dots, r_j$ for $x_1, \dots, x_j$ in p and indicating the remaining free variables in p. We use $\text{realRoots}(p) \subseteq \mathbb{R}^i$ to denote the set of *real roots of* p, $\deg_{x_j}(p)$ to denote the *degree of* p *in* x_j, $\text{coeff}_{x_j}(p)$ the set of *coefficients of* p *in* x_j, $\text{ldcf}_{x_j}(p)$ the *leading coefficient of* p *in* x_j, $\text{factors}(p)$ to denote the *irreducible factors of* p, $\text{disc}_{x_j}(p)$ to denote the *discriminant of* p *with respect to* x_j, and $\text{res}_{x_j}(p, q)$ to denote the *resultant of* p *and* q *with respect to* x_j. Let $r \in \mathbb{R}^{i-1}$ then p *is nullified on* r if $p(r, x_i) = 0$.

A *constraint* $p \sim 0$ compares a polynomial $p \in \mathbb{Q}[x_1, \dots, x_i]$ to zero using a relation symbol $\sim \in \{=, \neq, <, >, \leq, \geq\}$, and has the *solution set* $\{r \in \mathbb{R}^i \mid p(r) \sim 0\}$. A subset of $\mathbb{R}^i$ for some $i \in [n]$ is called *semi-algebraic* if it is the solution set of a Boolean combination of polynomial constraints. A *cell* is a non-empty connected subset of $\mathbb{R}^i$ for some $i \in [n]$. A polynomial $p \in \mathbb{Q}[x_1, \dots, x_i]$ is *sign-invariant on a set* $R \subseteq \mathbb{R}^i$ if the sign of $p(r)$ is the same for all $r \in R$.

Given $i, j \in \mathbb{N}_{>0}$ with $j < i$, we define the *projection of a set* $R \subseteq \mathbb{R}^i$ *onto* $\mathbb{R}^j$ by $R\downarrow_{[j]} = \{(r_1, \dots, r_j) \mid \exists r_{j+1}, \dots, r_i.\ (r_1, \dots, r_i) \in R\}$. Given a cell $R \subseteq \mathbb{R}^i$, $i \in [1..n]$ and continuous functions $f, g : R \to \mathbb{R}$, we define the cells $R \times (f, g) = \{(r, r_{i+1}) \mid r \in R, r_{i+1} \in (f(r), g(r))\}$ ($R \times (-\infty, g)$, $R \times (f, \infty)$ analogously).

An i-dimensional *(analytic) submanifold of* $\mathbb{R}^n$ is a non-empty subset $R \subseteq \mathbb{R}^n$ that "looks locally like $\mathbb{R}^i$". Given an open subset $U \subseteq \mathbb{R}^i$, a function $f : U \to \mathbb{R}$ is called *analytic* if it has a multiple power series representation around each point of U. Given an i-dimensional submanifold R of $\mathbb{R}^n$, a function $f : R \to \mathbb{R}$ is called *analytic* if for all $r \in R$, R looks locally like $\mathbb{R}^i$ with respect to a coordinate system about r and f looks locally like an analytic function $\mathbb{R}^i \to \mathbb{R}$. Let $p \in \mathbb{Q}[x_1, \dots, x_n]$ be a polynomial and $r \in \mathbb{R}^n$ be a point. Then the *order* $\text{ord}_r(p)$ *of* p *at* r is defined as the minimum k such that some partial derivative of total order k of p does not vanish at r (and ∞ if all vanish). We call p *order-invariant on* $R \subseteq \mathbb{R}^n$ if $\text{ord}_r(p) = \text{ord}_{r'}(p)$ for all $r, r' \in R$. For details, we refer to [8].

2.1 Cylindrical Algebraic Decomposition, McCallum's Projection and Levelwise Single Cell Construction

We give a short reminder of the *cylindrical algebraic decomposition (CAD)* [5] and the projection operator by McCallum [8, 9]: A CAD is a decomposition $\mathcal{C}$ of $\mathbb{R}^n$ such that each cell $R \in \mathcal{C}$ is semi-algebraic and *locally cylindrical* - i.e. can be described as solution set of $\psi_1(x_1) \wedge \psi_2(x_1, x_2) \wedge \psi_n(x_1, \dots, x_n)$ where ψ_i is either $x_i = \theta(x_1, \dots, x_{i-1})$, $\theta_l(x_1, \dots, x_{i-1}) < x_i < \theta_u(x_1, \dots, x_{i-1})$, $\theta_l(x_1, \dots, x_{i-1}) < x_i$, or $x_i < \theta_u(x_1, \dots, x_{i-1})$ for some continuous functions $\theta, \theta_l, \theta_u$ - and $\mathcal{C}$ is *cylindrically arranged* - i.e. either $n = 1$ or $\{R\downarrow_{n-1} \mid R \in \mathcal{C}\}$ is a cylindrically arranged decomposition of $\mathbb{R}^{n-1}$. The shape of such a CAD allows reasoning about properties of (sets of) polynomials computationally. In particular, it is called *sign-invariant for a set of polynomials* $P \subseteq \mathbb{Q}[x_1, \dots, x_n]$ if each $p \in P$ is sign-invariant on each $R \in \mathcal{C}$. A sign-invariant CAD for P is computed recursively, i.e. to describe the cells' boundaries for x_n, we first compute the underlying decomposition by a projection operation resulting in a set $P' \subseteq \mathbb{Q}[x_1,\dots,x_{n-1}]$ whose sign-invariant CAD will describe the first $n-1$ dimensions of the cells of the sign-invariant CAD of P. The *single cell construction* [4, 10] method computes, given a set of polynomials $P \subseteq \mathbb{Q}[x_1,\dots,x_n]$ and a sample point $s \in \mathbb{R}^n$, a locally cylindrical cell $R \subseteq \mathbb{R}^n$ such that $s \in \mathbb{R}^n$ and such that P is sign-invariant on R. In the rest of this section, we give a brief introduction to the *levelwise method* [10] and the required CAD theory.

Delineability. A central notion states that the variety of a polynomial can be described by continuous functions which are nicely ordered over a given cell. This allows us to reason about the polynomial's roots using these functions.

Definition 1 (Delineability [5, 9]). *Let $i \in \mathbb{N}$, $R \subseteq \mathbb{R}^i$ be a cell, and $p \in \mathbb{Q}[x_1,\dots,x_{i+1}] \setminus \{0\}$. The polynomial p is called* delineable on R *if and only if there exist finitely many continuous functions $\theta_1,\dots,\theta_k : R \to \mathbb{R}$ (for $k \geq 0$) such that*

- *$\theta_1 < \ldots < \theta_k$;*
- *the set of real roots of $p(r, x_{i+1})$ is $\{\theta_1(r),\dots,\theta_k(r)\}$ for all $r \in R$; and*
- *there exist constants $m_1,\dots,m_k \in \mathbb{N}_{>0}$ such that for all $r \in R$ and all $j \in [1..k]$, the multiplicity of the root $\theta_j(r)$ of $p(r, x_{i+1})$ is m_j.*

The θ_j are called real root functions of p on R. *The sets $R \times \theta_j$ are called* sections of p over R. *The cells in $R \times \mathbb{R}$ minus these sections are called* sectors of p over R.

Analytic delineability *is defined like delineability, but the underlying cell is required to be a connected analytic submanifold and the real root functions are required to be analytic.* △

In particular, if a polynomial is delineable on some cell, then we can refer to each root function by an index. Let $i \in \mathbb{N}$, $p \in \mathbb{Q}[x_1,\dots,x_{i+1}]$, $\mathrm{level}(p) = i+1$, and $j \in \mathbb{N}_{>0}$. An *indexed root* is a partial function $\mathrm{root}_{x_{i+1}}[p,j] : \mathbb{R}^i \hookrightarrow \mathbb{R}$ that maps $s \in \mathbb{R}^i$ to the j-th real root of $p(s, x_{i+1})$ if it exists. Given a cell $R \subseteq \mathbb{R}^i$ where p is delineable, then $\mathrm{root}_{x_{i+1}}[p,j]$ coincides with the root function θ_j from the above definition on R; to simplify notation, we use both notions interchangeably when R is clear from the context. Let θ denote the above indexed root, then $\theta.p$ and $\theta.j$ refer to p and j respectively. We denote the set of indexed roots of p that are defined at $s \in \mathbb{R}^i$ by $\mathrm{irExpr}(p, s)$; we define this set analogously for sets of polynomials as well.

The following theorem gives a projection to obtain a cell where a polynomial is delineable:

Theorem 1 (Delineability of a Single Polynomial [9, Theorem 2], [3, Theorem 3.1]). *Let $i \in \mathbb{N}$, $R \subseteq \mathbb{R}^i$ be a connected analytic submanifold, and $p \in \mathbb{Q}[x_1,\dots,x_{i+1}]$, $\mathrm{level}(p) = i+1$ be irreducible. Assume that p is not nullified at any point in R, $\mathrm{disc}_{x_{i+1}}(p)$ is order-invariant on R, and $\mathrm{ldcf}_{x_{i+1}}(p)$ is sign-invariant on R.*

Then p is analytically delineable on R and is order-invariant on its sections over R. △

Root Orderings. Once we can describe the roots of individual polynomials by ordered root functions on the underlying cell, we can reason about intersections of root functions from different polynomials, e.g. ensure that two root functions remain in the same order on the underlying cell.

Theorem 2 (Lifting of Pairs of Polynomials [10, Theorem A.1]). *Let $i \in \mathbb{N}$, $R \subseteq \mathbb{R}^i$ be a connected analytic submanifold, $s \in R$, and $p_1, p_2 \in \mathbb{Q}[x_1, \ldots, x_{i+1}]$ be irreducible and coprime such that* $\text{level}(p_1) = \text{level}(p_2) = i+1$. *Assume that p_1 and p_2 are analytically delineable on R and $\text{res}_{x_{i+1}}(p_1, p_2)$ is order-invariant on R.*

Let $\theta_1, \theta_2 : R \to \mathbb{R}$ be real root functions of p_1 and p_2 on R respectively, and $\sim \in \{<, =\}$ such that $\theta_1(s) \sim \theta_2(s)$. Then $\theta_1 \sim \theta_2$ on R. △

In fact, to maintain that two real root functions θ_1 and θ_2 stay in the same order on R, we could also exploit transitivity, e.g. $\theta_1 < \theta_3$ on R and $\theta_3 < \theta_2$ on R implies $\theta_1 < \theta_2$ on R. The work in [10] generalizes this idea to orderings in a set of root functions. This allows for flexibility in the choice of resultants which we compute to maintain certain invariance properties, potentially avoiding the computation of expensive resultants. For this paper, we do not consider exploiting transitivity, although everything can be extended accordingly.

Single Cell Construction. The idea is, given polynomials $P \subseteq \mathbb{Q}[x_1, \ldots, x_n]$ and a sample $s \in \mathbb{R}^i$, we compute and sort the real roots of $p(s_{[i-1]}, x_i),\ p \in P$. We determine the greatest root below (or equal to) s_i and the smallest root above (or equal to) s_i (if they do not exist, we use $-\infty$ and ∞ respectively). Let θ_l and θ_u be the corresponding real root functions respectively. To describe a sign-invariant cell for P, the bounds on x_i are described by the *symbolic interval (of level i)* (θ_l, θ_u) (whose bounds depend on $x_1, \ldots, x_{[i-1]}$) if $\theta_l(s_{[i-1]}) < \theta_u(s_{[i-1]})$ and $[\theta_l, \theta_u]$ if $\theta_l(s_{[i-1]}) = \theta_u(s_{[i-1]})$. We use $\mathtt{I}.l$ and $\mathtt{I}.u$ to refer to θ_l and θ_u respectively. We define $R \times \mathtt{I}$ as $R \times (\theta_l, \theta_u)$ (first case) or $R \times \theta_l.l$ (second case). The idea is now to use root orderings to make sure that each $p \in P$ is sign-invariant in the interval:

Theorem 3 (Root Ordering for Sign Invariance). *Let $i \in \mathbb{N}$, $R \subseteq \mathbb{R}^i$ be connected, $p \in \mathbb{Q}[x_1, \ldots, x_{i+1}]$ be irreducible such that* $\text{level}(p) = i + 1$, $\mathtt{I}$ *be a symbolic interval of level $i + 1$. Assume that p, $\mathtt{I}.l.p$, $\mathtt{I}.u.p$ (if defined) are analytically delineable on R, and that for each real root function θ of p on R it holds*

- *if $\mathtt{I} = (\mathtt{I}.l, \mathtt{I}.u)$, then either $\mathtt{I}.l \neq -\infty$ and $\theta \sim \mathtt{I}.l$ on R for some $\sim \in \{<, =\}$, or $\mathtt{I}.u \neq -\infty$ and $\mathtt{I}.u \sim \theta$ on R for some $\sim \in \{<, =\}$;*
- *if $\mathtt{I} = [\mathtt{I}.l, \mathtt{I}.u]$ with $\mathtt{I}.l = \mathtt{I}.u$, then either $\theta < \mathtt{I}.l$ on R, $\mathtt{I}.u < \theta$ on R, or $\theta = \mathtt{I}.l = \mathtt{I}.u$ on R.*

Then p is sign-invariant on $R \times \mathtt{I}$. △

We can compute witnesses for the real root functions of p on R by computing the set $\text{irExpr}(p, s_{[i-1]})$. This set covers all roots of p that might appear on R if p is delineable on R. This results in the abstract algorithm in Algorithm 1.

Algorithm 1: `single_cell_construction(`P`,`s`)`

Input : finite $P \subset \mathbb{Q}[x_1, \ldots, x_n]$, $s \in \mathbb{R}^n$
Output: Symbolic intervals $\mathtt{I}_1, \ldots, \mathtt{I}_n$ for $x_1, \ldots, x_n$ describing a sign-invariant cell for P containing s

1 **foreach** $i = n, \ldots, 1$ **do**
2 $\quad P_i := \{p \in P \mid \text{level}(p) = i\}$, $P := P \setminus P_i$
3 $\quad \{\theta_1, \ldots, \theta_k\} := \text{irExpr}(P_i, s_{[i-1]})$ such that $\theta_1(s_{[i-1]}) \leq \ldots \leq \theta_k(s_{[i-1]})$
$\quad$ `// Determine symbolic interval`
4 $\quad$ **if** $s_i = \theta_j(s_{[i-1]})$ for some j **then** $\mathtt{I}_i := [\theta_j, \theta_j]$
5 $\quad$ **else if** $\theta_j(s_{[i-1]}) < s_i < \theta_{j+1}(s_{[i-1]})$ for some j **then** $\mathtt{I}_i := (\theta_j, \theta_{j+1})$
6 $\quad$ **else if** $s_i < \theta_1(s_{[i-1]})$ for some j **then** $\mathtt{I}_i := (-\infty, \theta_1)$
7 $\quad$ **else if** $\theta_k(s_{[i-1]}) < s_i$ for some j **then** $\mathtt{I}_i := (\theta_k, \infty)$
8 $\quad$ **else** $\mathtt{I}_i := (-\infty, \infty)$
9 $\quad$ **foreach** $p \in P_i$ **do**
$\qquad$ `// Ensure order invariance for each polynomial`
10 $\qquad$ **if** $p(s_{[i-1]}, x_i) = 0$ **then return** FAIL
11 $\qquad$ **add** some $c \in \text{coeff}_{x_i}(p)$ to P such that $c(s) \neq 0$
12 $\qquad$ **add** polynomials to P s.t. p is delineable and order-invariant in its sections according to Theorem 1
13 $\qquad$ **add** polynomials to P s.t. p is sign-inv. acc. to Theorems 2 and 3
14 $\quad$ **add** polynomials to P s.t. $\mathtt{I}_i.l < \mathtt{I}_i.u$ according to Theorem 2
15 **return** $\mathtt{I}_1, \ldots, \mathtt{I}_n$

Note that in Line 10, the algorithm might fail, as McCallum's projection operator cannot reason about cells where a polynomial is nullified. In Line 11, we add additional coefficients that prevent polynomials from nullification on the constructed underlying cell, for which we refer to [10]. Further, we recall that the constructed call is an analytic submanifold as it is bounded by root functions. Line 13 ensures the connectedness of the cell.

3 Detecting "Irrelevant" Roots

Our main contribution in this work is to merge adjacent cells during single cell construction by ignoring root functions when determining the symbolic interval describing the bounds on some variable. For that we identify which roots are "irrelevant". Before we formally define what we mean by this, we illustrate the ideas by some examples.

Example 1. Consider the polynomials $p_1 = x_2^2 + x_3^2 - 1$ and $p_2 = -1.25 + x_2 + x_3$ and the sample $s = (0, -1.125)$, whose real roots are depicted in Figure 1 (note that we do not use x_1 here yet, the example is two-dimensional). The single cell algorithm would produce a cell where the x_3-dimension is bounded from above by the *first* root of p_1. We make sure it is well-defined by adding the discriminant p_4 of p_1. To maintain that no root of p_2 crosses this upper bound over the underlying cell, we add the resultant p_5 of p_1 and p_2 to the projection; speaking in terms of properties, we maintain the ordering $\text{root}_{x_2}[p_1, 1] < \text{root}_{x_2}[p_2, 1]$ on

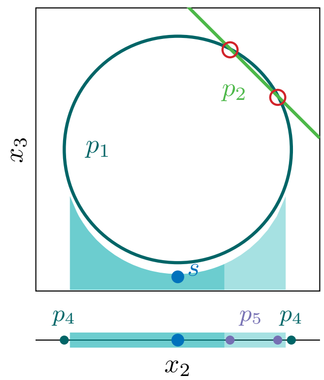

Fig. 1: First example. The intersections of the real roots of p_1 and p_2 (marked with red circles) are irrelevant for constructing the cell around s. Thus, instead of the smaller cell (shaded in darker colour) bounded from above by a root of the resultant p_5 in the x_2-dimension, we could obtain the wider cell (shaded in lighter colour) bounded from above by the discriminant p_4 in the x_2-dimension.

the one-dimensional cell. Now, for the x_2-dimension, the cell will be bounded from above by a root of p_5. However, both roots of the resultant mark points where p_2 intersects the *second* root of p_1, thus these intersections would not affect sign-invariance of the constructed cell - they are *irrelevant* to the constructed cell. In this example, we thus could ignore the roots of p_5, yielding a much larger cell. Computationally, we need to check for every root $s_2 \in \mathbb{R}$ of p_5 whether the first root of $p_1(s_2, x_3)$ is smaller than the first root of $p_2(s_2, x_3)$ - the computations refer to the correct roots as p_2 is delineable on $\mathbb{R}$ and p_1 is delineable on an interval containing the sample point and the two roots of the resultant (and thus the roots can be indexed consistently over that interval).

In fact, to describe the same set of points, the single cell would need to compute 5 cells (all upper bounded by the first root of p_1 in the x_3-dimension, and in the x_2-dimension bounded by the point intervals defined by the roots of p_5, and the three open intervals between the roots of p_4 and p_5 respectively). Thus, by ignoring roots of the resultant, we *merge* adjacent cells.

Now consider the modified polynomials $p_1 = x_2^2 + x_3^2 - 1$ and $p_2 = x_1 + x_2 + x_3$ and the sample $s = (-1.25, 0, -1.125)$. The circle p_1 becomes a cylinder aligned with the x_1-axis, and the line gets a plane which is not parallel to any axis. Five cuts of this example and its projection are depicted in Figure 2.

Consider the third cut at $x_1 = -1.25$, where the sample point lies. In fact, the situation is the same as in Figure 1, thus we can determine irrelevant intersections analogously to obtain a larger cell. However, note that we cannot ignore the corresponding roots of the resultant p_5 completely: consider the cut at 0, where the intersection became relevant, i.e. if we would not bound the cell in the x_1-dimension properly, then the cell is not sign-invariant any more. The intuitive reason is that at $x_1 = -1$, the second root of the resultant p_5 switches from an intersection of p_2 with the *second* root of p_1 to an intersection with the *first* root of p_1, which describes the upper bound in the x_3-dimension. This point is

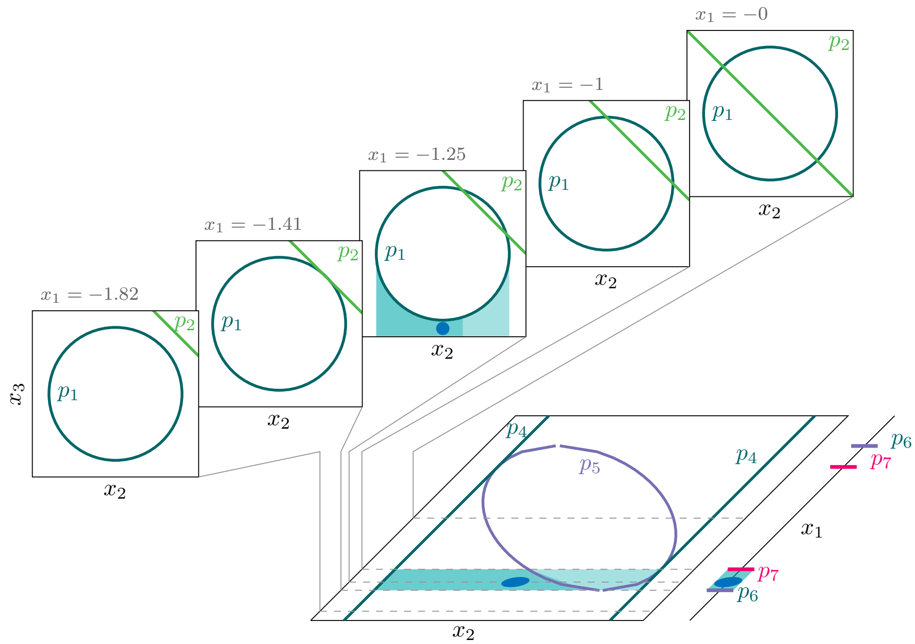

Fig. 2: First example, extended by another dimension. The five coordinate systems at the top left depict x_1-cuts. The two-dimensional coordinate system at the bottom depicts the CAD-projection in the x_1-x_2-plane, and the one-dimensional coordinate system depicts the CAD projection on the x_1-line. On all coordinate systems: the smaller cell is shaded in darker colour, the larger cell is shaded in lighter colour, the point is the sample.

caught by the intersection of the resultant p_5 with the discriminant p_4 of p_1, i.e. we need to bound the cell in x_1-dimension by the root of the resultant p_7 of p_4 and p_5. For similar reasons, we need to take the discriminant p_6 of the resultant p_5 into account in general.

The reason why we catch everything "bad" that can happen is: The sections of the resultant p_5 of p_1 and p_2 are well-ordered between the roots of the discriminant p_4 of p_1 - i.e. the variety of p_5 and p_4 can be described by continuous functions $\theta_1 < \theta_2 < \theta_3 < \theta_4$ on the interval between the corresponding roots of p_6 and p_7. Thus, the sections θ_2 and θ_3 mark the points where the *same* roots of p_1 and p_2 intersect. These intersections are relevant at the current sample point (which we can check computationally as described above) if and only if these intersections are relevant everywhere above these sections. We will later define a property called *local delineability* that generalizes the well-ordering of the sections of the resultants and discriminants. △

Example 2. Consider the same set of polynomials, now with the sample $s = (0, 0, -1.125)$.

As depicted in Figure 3, only a single intersection of p_1 and p_2 is irrelevant to the cell (we maintain the ordering $\text{root}_{x_2}[p_1, 1] < \text{root}_{x_2}[p_2, 1]$). Still, we can

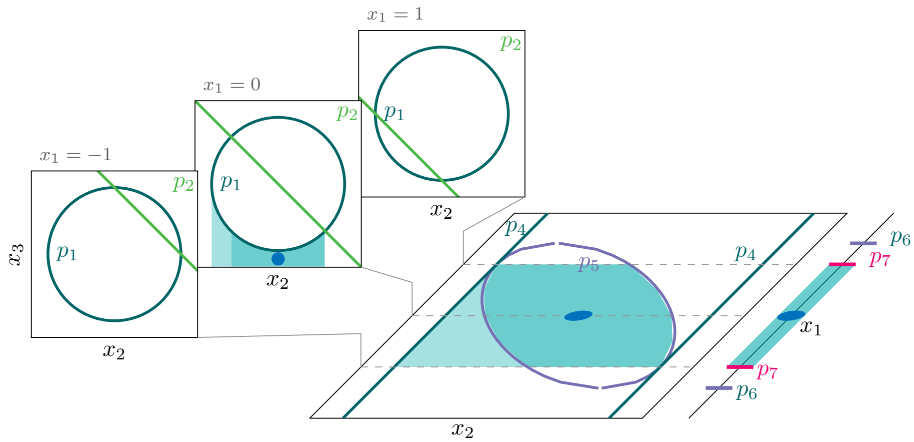

Fig. 3: First example, different sample point.

merge three adjacent cells (which are in the same cylinder). As the second root of the resultant p_5 witnesses the non-irrelevant intersection, it defines the upper bound in the x_2-dimension of the cell. △

Before we formally define what local delineability is, we examine a second example hinting an even more general observation:

Example 3. Consider the sample point $(-0.5, 0, -1.125)$, the set of polynomials $\{p_1, p_2, p_3\}$ and its projection polynomials:

$$
\begin{array}{lll}
p_1 = x_2^2 + x_3^2 - 1 & p_2 = x_3 + 0.75 & p_3 = x_1 - x_3 \\
p_4 = \mathrm{disc}_{x_3}(p_1) & p_5 = \mathrm{res}_{x_3}(p_1, p_3) & p_6 = \mathrm{res}_{x_3}(p_1, p_2) \\
p_7 = \mathrm{disc}_{x_2}(p_5) & p_8 = \mathrm{res}_{x_2}(p_5, p_6) & p_9 = \mathrm{res}_{x_2}(p_6, p_4)
\end{array}
$$

As in the previous examples, we depict some cuts and the projection in Figure 4. Observe that the original cell is bounded by the iterated resultant p_8, i.e. over its roots, the intersections of p_3 with p_1 and p_2 with p_1 intersect. For the cell to be constructed around the given sample point, this intersection is not relevant to our cell[1], as p_3 intersects with the *second* root of p_1, while the *first* root of p_1 defines the upper bound in the x_3-dimension (that is, we maintain the ordering $\mathrm{root}_{x_2}[p_1, 1] < \mathrm{root}_{x_2}[p_2, 1]$ and $\mathrm{root}_{x_2}[p_1, 1] < \mathrm{root}_{x_2}[p_3, 1]$ on the two-dimensional cell). We can thus extend the cell beyond the root of p_8.

Now observe that the variety of the resultant p_5 of p_1 and p_2 and the discriminant p_4 of p_1 in the larger underlying two-dimensional cell is not described by well-ordered continuous functions - but it is over some larger cell (in the same cylinder over the one-dimensional cell). In this case, to maintain the mentioned *local delineability* property, we only need the discriminant p_7 of p_5, the resultant of p_6 and p_4 (which is trivial), and the resultant p_9 of p_5 with p_4 - we do not even need to compute the iterated resultant p_8 at all! In contrast to the previous example, the number of computed projection polynomials is reduced. △

[1] If we would construct a cell at $(\mathbf{0.5}, 0, -1.125)$, this intersection of intersections *would be* relevant to the cell.

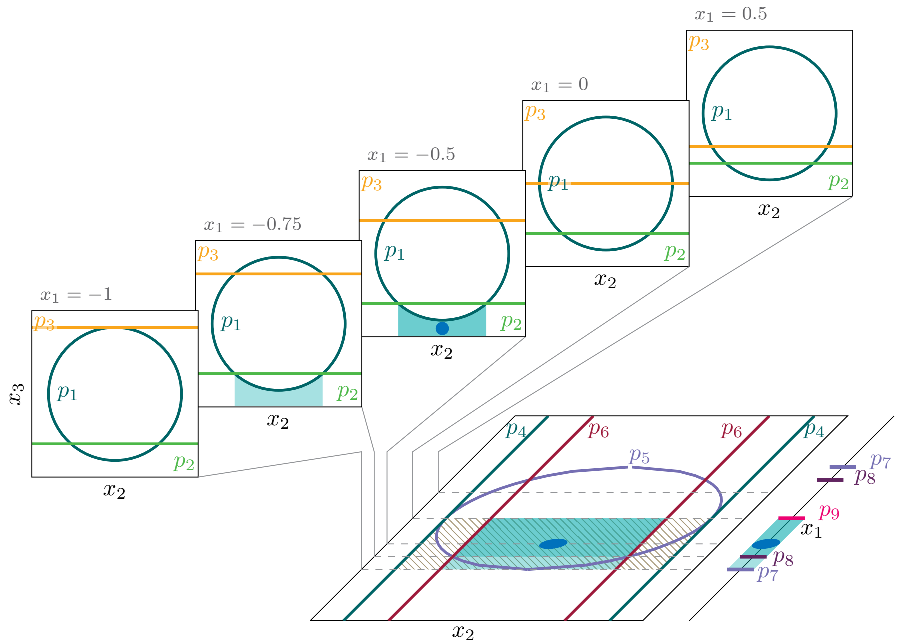

Fig. 4: Second example. In the hatched cell, the roots of p_4 and p_5 are well-ordered.

4 Local Delineability and Modified Projection

In the previous section, we developed ideas for an adaption of Algorithm 1 for detecting and generalizing irrelevant roots. We will extend the set P of projection polynomials for pairs of indexed roots whose order needs to be maintained; that is, instead of adding projection polynomials directly in Line 13, we add pairs of root functions that ensure order-invariance of the given polynomial. Analogously to polynomials, pairs of root functions are processed using theorems analogous to Theorems 1 and 2 which we introduce in this section.

Processing a pair of root functions involves the following steps: (1) We check over which roots of the resultant of the defining polynomials the root functions intersect. (2) We ensure that the respective resultant is *locally delineable* in the underlying cell, that is, its root functions do not intersect in the underlying cell. This establishes a correspondence between the intersections of the root functions and the resultant's sections. (3) We merge the adjacent sectors with a resultant's section when the root functions do not intersect above the section (i.e. the test from step (1) fails).

To elaborate step (1): Assume two root functions θ_1 and θ_2 of two polynomials p_1 and p_2, both of level $i+1$. Remind their representation as indexed root expressions: If p_1 and p_2 are delineable over some cell, the index of their roots is invariant, i.e. the indexed root will refer to the same root function. We thus determine the interval I such that p_1 and p_2 are delineable on $s_{[i-1]} \times I$ by computing the required coefficients and discriminants. We now consider each root $s' \in \mathbb{R}$ of the resultant $\mathrm{res}_{x_i}(p_1, p_2)$ over the current sample $s_{[i-1]} \in \mathbb{R}^{i-1}$ to

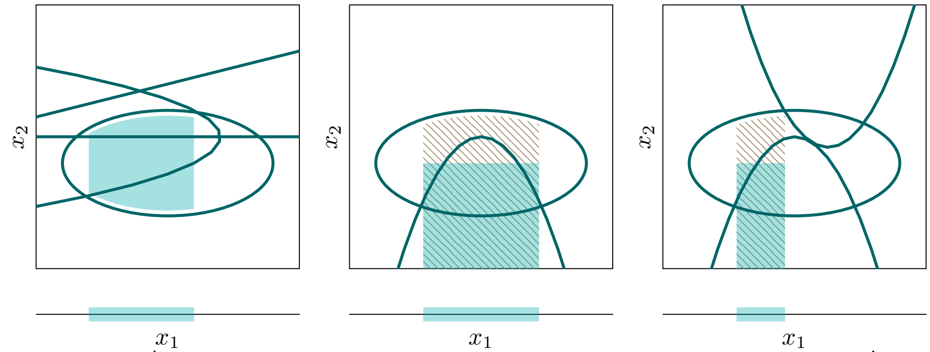

(a) Within R', no intersections of the roots are allowed; above or below R', intersections are allowed.

(b) It is also possible that roots "leave" and re-enter R', if all roots are well-ordered in R''.

(c) If a root leaves R' and re-enters R', no intersections with another root between these points are allowed.

Fig. 5: Examples of local delineability. The shaded area in the one-dimensional coordinate system depicts R from Definition 2, the shaded area in the two-dimensional one depicts R', the hatched area depicts R'' if it differs from R'.

determine whether θ_1 and θ_2 intersect above $(s_{[i-1]}, s')$: If $s' \in I$, we evaluate θ_1 and θ_2 at $(s_{[i-1]}, s')$ and check whether the result is equal. If $s' \notin I$, we cannot check for intersections of the root functions, thus we assume s' is relevant. We will formalize this in Theorem 5.

In step (2), we will determine an *ordering of the resultant's roots* that needs to hold on the underlying cell $R' \subseteq \mathbb{R}^{i-1}$ such that the resultant is locally delineable on a part of the cylinder $R' \times \mathbb{R}$ around the current sample point.

In step (3), we determine the set R: It will be composed of some adjacent sections and sectors $R_1, \ldots, R_k \subseteq \mathbb{R}^i$ on R' of the resultant such that θ_1 and θ_2 are in the same order on $R_1, \ldots, R_k$ (and thus on R). To determine these sections and sectors, we use the information computed in step (1); the local delineability in step (2) ensures that these cells are *arranged cylindrically*, i.e. the cells that we merge are adjacent cells in the same cylinder $R' \times \mathbb{R}$. Thus, we obtain a cylindrical cell R after the merge.

We formalize step (2):

Definition 2 (Local delineability). *Let* $i \in \mathbb{N}_{>0}$*,* $R \subseteq \mathbb{R}^i$ *be a cell,* $p \in \mathbb{Q}[x_1, \ldots, x_{i+1}]$*,* $\text{level}(p) = i+1$*, and* $R' \subseteq \mathbb{R}^{i+1}$ *connected such that* $R' \downarrow_{[i]} = R$*. The polynomial* p *is called* locally R'-delineable *if and only if there exists a cell* $R'' \supseteq R'$ *such that* $R'' \downarrow_{[i]} = R$ *and there exist finitely many continuous functions* $\theta_1, \ldots, \theta_k : R \to \mathbb{R}$ *(for* $k \leq 0$*) such that*

- $(r, \theta_j(r)) \in R''$ *for all* $r \in R$ *and all* $j \in [1..k]$*;*
- $\theta_j < \theta_{j+1}$ *on* R *for all* $j \in [1..k-1]$*;*
- *the set of real roots of the univariate polynomial* $p(r, x_{i+1})$ *that are contained in* R'' *is* $\{\theta_j(r) \mid i \in [1..k], r \in R\}$ *for all* $r \in R$*; and*
- *there exist constants* $m_1, \ldots, m_k \in \mathbb{N}_{>0}$ *such that for all* $j \in [1..k]$ *and all* $r \in R$*, the multiplicity of the root* $\theta_j(r)$ *of* $p(r, x_{i+1})$ *is* m_j*.*

A set of polynomials $P \subset \mathbb{Q}[x_1, \ldots, x_{i+1}]$*,* $\mathrm{level}(\prod_{p \in P} p) = i + 1$ *is called* locally R'-delineable *if and only if* $\prod_{p \in P | \mathrm{level}(p)=i+1} p$ *is* locally R'-delineable*.*

Analytic local delineability *is defined as local delineability, with the modification that it is only defined on connected analytic submanifolds (instead of general cells) and the functions* $\theta_1, \ldots, \theta_k$ *are required to be analytic (instead of only continuous).* △

Note that if a polynomial is delineable on R, then it is locally $R \times (-\infty, \infty)$-delineable. Figure 5 illustrates various examples of local delineability.

We now give a theorem that allows to realize local delineability in an algorithm, assuming we can maintain an ordering of pairs of root functions (as in Theorem 2 or the later following Theorem 5):

Theorem 4 (Local Delineability of a Set of Polynomials). *Let* $i \in \mathbb{N}$*,* $R \subseteq \mathbb{R}^i$ *be a connected analytic submanifold,* $s \in R$*,* $P \subset \mathbb{Q}[x_1, \ldots, x_{i+1}]$ *such that* $\mathrm{level}(\prod_{p \in P} p) = i + 1$*. Let* $\mathtt{I}$ *be a symbolic interval of level* $i + 1$*, and let* $\mathrm{irExpr}(\mathrm{factors}(P), s) = \{\theta_1, \ldots, \theta_k\}$ *s.t.* $\theta_1(s) \leq \ldots \leq \theta_k(s)$*.*

Assume that $\mathtt{I}.l.p$*,* $\mathtt{I}.u.p$ *are analytically delineable on* R*, each* $q \in \mathrm{factors}(P)$ *with* $\mathrm{level}(q) < i + 1$ *is order-invariant on* R*, and for each* $q \in \mathrm{factors}(P)$ *with* $\mathrm{level}(q) = i + 1$ *it holds*

- q *is not nullified on any point in* R*,*
- $\mathrm{disc}_{x_{i+1}}(q)$ *is order-invariant on* R*,*
- $\mathrm{ldcf}_{x_{i+1}}(q)$ *is sign-invariant on* R*.*

Assume that there exist $\ell, \ell' \in [1..k]$ *such that:*

1. *For all* $j \in [\ell..\ell'{-}1]$ *it holds* $\theta_j \sim \theta_{j+1}$ *on* R *for some* $\sim \in \{<, =\}$*.*
2. *For all* $j \in [1..\ell{-}1]$ *it holds*
   - *if* $\mathtt{I} = (\mathtt{I}.l, \mathtt{I}.u)$ *with* $\mathtt{I}.l \neq -\infty$*, then* $\theta_j \sim \mathtt{I}.l$ *on* R *for some* $\sim \in \{<, =\}$*;*
   - *if* $\mathtt{I} = [\mathtt{I}.l, \mathtt{I}.u]$*, then* $\theta_j < \mathtt{I}.l$ *on* R*.*
3. *For all* $j \in [\ell'{+}1..k]$ *it holds*
   - *if* $\mathtt{I} = (\mathtt{I}.l, \mathtt{I}.u)$ *with* $\mathtt{I}.u \neq \infty$*, then* $\mathtt{I}.u \sim \theta_j$ *on* R *for some* $\sim \in \{<, =\}$*;*
   - *if* $\mathtt{I} = [\mathtt{I}.l, \mathtt{I}.u]$*, then* $\mathtt{I}.u < \theta_j$ *on* R*.*
4. *For all* $j \in [1..\ell{-}1]$ *it holds* $\theta_j \sim \theta_\ell$ *on* R *for some* $\sim \in \{<, =\}$*.*
5. *For all* $j \in [\ell'{+}1..k]$ *it holds* $\theta_{\ell'} \sim \theta_j$ *on* R *for some* $\sim \in \{<, =\}$*.*

Then P *is analytically locally* $(R \times \mathtt{I})$*-delineable, and* $q \in \mathrm{factors}(P)$ *is order-invariant in each section of* q *on* R *that have a non-empty intersection with* $(R \times \mathtt{I})$*.* △

Note that if we choose $\ell = 1$ and $\ell' = k$, then this yields delineability of P on R (see Example 3). If $\ell' < \ell$, then this yields sign-invariance of P on R.

Proof. Assume that R fulfils the antecedents of the theorem.

Let $\theta'_1, \ldots, \theta'_{k'} : R \to \mathbb{R}$ be the real root functions describing the graph of $\prod_{p \in P} p$ on R. By construction, for every $j' \in [1..k']$ there exists $j \in [1..k]$ such that $\theta_j = \theta'_{j'}$ on R. Note that the lower level factors are not relevant for local delineability.

By analysis of each case, the real root functions of P defined at s are in the shape required by Definition 2: $\theta_\ell, \dots, \theta_{\ell'}$ are the root functions matching Definition 2 contained in some $R'' \subseteq \mathbb{R}^{i+1}$ such that $R'' \downarrow_{[i]} = R$. Condition 1 ensures that they do not intersect. θ_ℓ and $\theta_{\ell'}$ are the functions that bound R''; Conditions 4 and 5 ensure that all other root functions $\theta_1, \dots, \theta_{\ell-1}, \theta_{\ell'+1}, \dots, \theta_k$ ($\theta_1, \dots, \theta_k$ if $l > l'$) do not enter R''. $R' = (R \times \mathtt{I})$ is described by $\mathtt{I}$; Condition 2 and 3 ensure that $R' \subseteq R''$. Thus, the requirements from Definition 2 are fulfilled, and P is analytically locally $(R \times \mathtt{I})$-delineable.

The lower-level factors are all order-invariant on R by assumption, and the factors on the current level are order-invariant in each of their respective sections on R by Theorem 1. □

The above theorem only gives the conditions that need to hold for local delineability. Now, we formalize which roots are irrelevant, and what needs to hold in the underlying projection:

In the examples above, we hinted that we can only distinguish relevant from irrelevant intersections of root functions where the defining polynomials are delineable, as we can only compute values of indexed roots. Thus, for $i \in \mathbb{N}$, $p \in \mathbb{Q}[x_1, \dots, x_{i+1}]$, and $s \in \mathbb{R}^i$, let $\mathrm{proj\_polys}(p)$ be the projection polynomials of level i that need to be sign-invariant to maintain delineability of p according to Theorem 1 (i.e. the discriminant, the leading coefficient, and some non-zero coefficient for non-nullification). Further, for $P \subset \mathbb{Q}[x_1, \dots, x_{i+1}]$, let $\rho_l = \max\{\rho \in \mathrm{realRoots}(\cup_{p \in P, q \in \mathrm{proj\_polys}(p)} q(s_{[i-1]})) \mid \rho \leq s_i\}$ and $\rho_u = \min\{\rho \in \mathrm{realRoots}(\cup_{p \in P, q \in \mathrm{proj\_polys}(p)} q(s_{[i-1]})) \mid \rho \geq s_i\}$; we define the *delineable interval* of P w.r.t. s as $\mathrm{del\_int}(P, s) = (\rho_l, \rho_u)$ if $\rho_l \neq \rho_u$ and $\mathrm{del\_int}(P, s) = [\rho_l, \rho_u]$ otherwise.

Theorem 5 (Root Orderings for Pairs of Root Functions / Lifting of Pairs of Root Functions). *Let $i \in \mathbb{N}$, $R \subseteq \mathbb{R}^i$ be a connected analytic submanifold, $s \in R$, $p_1, p_2 \in \mathbb{Q}[x_1, \dots, x_{i+1}]$ be irreducible and coprime such that $\mathrm{level}(p_1) = \mathrm{level}(p_2) = i + 1$, $\theta_1, \theta_2 : R \to \mathbb{R}$ be real root functions of p_1 and p_2 respectively such that $\theta_1(s) < \theta_2(s)$, $\mathtt{I}$ be a symbolic interval of level i, and $\mathrm{irExpr}(\mathrm{factors}(\mathrm{res}_{x_{i+1}}(p_1, p_2)), s_{[i-1]}) = \{\xi_1, \dots, \xi_k\}$ s.t. $\xi_1(s_{[i-1]}) \leq \dots \leq \xi_k(s_{[i-1]})$.*

Assume that p_1 and p_2 are analytically delineable on R, $\mathtt{I}.l.p$ and $\mathtt{I}.u.p$ are analytically delineable on $R \downarrow_{[i-1]}$, $R = R \downarrow_{[i-1]} \times \mathtt{I}$, $\{\mathrm{res}_{x_{i+1}}(p_1, p_2), proj\_polys(p_1), proj\_polys(p_2)\}$ is locally R-delineable and each of its factors q is order-invariant in each section of q on $R \downarrow_{[i-1]}$ that have a non-empty intersection with R, and for every $j \in [1..k]$, if

- *$\xi_j(s_{[i-1]}) \notin del\_int(\{p_1, p_2\}, s)$, or*
- *$s' := \xi_j(s_{[i-1]}) \in del\_int(\{p_1, p_2\}, s)$ and $\theta_1(s_{[i-1]}, s') = \theta_2(s_{[i-1]}, s')$,*

then it holds

- *if $\mathtt{I} = (\mathtt{I}.l, \mathtt{I}.u)$, then $\mathtt{I}.l \neq -\infty$ and $\xi_j \sim \mathtt{I}.l$ on $R \downarrow_{[i-1]}$ for some $\sim \in \{<, =\}$ or $\mathtt{I}.u \neq \infty$ and $\mathtt{I}.u \sim \xi_j$ on $R \downarrow_{[i-1]}$ for some $\sim \in \{<, =\}$;*

– *if* $\mathtt{I} = [\mathtt{I}.l, \mathtt{I}.u]$, *then* $\xi_j < \mathtt{I}.l$ *on* $R\downarrow_{[i-1]}$ *or* $\mathtt{I}.u < \xi_j$ *on* $R\downarrow_{[i-1]}$.

Then $\theta_1 < \theta_2$ *holds on* R. △

The adaption for the case $\theta_1(s) = \theta_2(s)$ is a degenerate case (the resultant is zero at s due to the intersection) and thus trivial. Note that checking the second bullet point requires lifting over each (possibly non-rational) root of the resultant, which might be computationally expensive. We emphasize that in an actual implementation, no complete check needs to be implemented. We could e.g. restrict ourselves to rational sample points.

Proof. Assume that R fulfils the antecedents of the theorem.

Assume for contradiction that there exists $t \in R$ such that $\theta_1(t) \not< \theta_2(t)$. As $\theta_1(s) < \theta_2(s)$, there exists some $t' \in R$ such that $\theta_1(t') = \theta_2(t')$ which is a zero of $\mathrm{res}_{x_{i+1}}(p_1, p_2)$, thus there exists $j^* \in [1..k]$ s.t. $\xi_{j^*}(t'_{[i-1]}) = t'_i$ (i.e. $t' \in R\downarrow_{[i-1]} \times \xi_{j^*}$).

Let $R' \supseteq R$, $R'\downarrow_{[i-1]} = R\downarrow_{[i-1]}$ be maximal such that $P = \{\mathrm{res}_{x_{i+1}}(p_1, p_2), \mathrm{proj\_polys}(p_1), \mathrm{proj\_polys}(p_2)\}$ is analytically R'-locally delineable on $R\downarrow_{[i-1]}$, and p_1 and p_2 are analytically delineable on R', and P is order-invariant in each of its sections on $R\downarrow_{[i-1]}$ that intersect with R'. Note that as we included the proj_polyssets into the set of polynomials to be locally delineable, for each $j \in [1..k]$ it holds that $(R\downarrow_{[i-1]} \times \xi_j) \cap R \neq \emptyset$ implies $R\downarrow_{[i-1]} \times \xi_j \subseteq R'$.

Now let $R'_1, \ldots, R'_{k'} \subseteq R'$ be the maximal sign-invariant cells of $\mathrm{res}_{x_{i+1}}(p_1, p_2)$ in R' (which are arranged cylindrically, that is for all $j \in [1..k']$ it holds $R'_j\downarrow_{[i-1]} = R\downarrow_{[i-1]}$); $\mathrm{res}_{x_{i+1}}(p_1, p_2)$ is order-invariant in each of these cells by assumption. By Theorem 2, for each $j \in [1..k']$ there exists $\sim_j$ such that $\theta_1 \sim_j \theta_2$ on R'_j.

Towards the contradiction, we conclude that ξ_{j^*} is one of the sections contained in R' (i.e. there exists $j \in [1..k']$ such that $R\downarrow_{[i-1]} \times \xi_{j^*} = R'_j$). By construction, it holds $(\{s_{[i-1]}\} \times \mathrm{del\_int}(\{p_1, p_2\}, s)) \cap R'_j = \{(s_{[i-1]}, s')\}$ for some $s' \in R'$ such that $s' = \xi_{j^*}(s_{[i-1]})$. By requirement of the theorem, ξ_{j^*} is required to “remain outside” I on $R\downarrow_{[i-1]}$, that is $(R\downarrow_{[i-1]} \times \xi_{j^*}) \cap R = \emptyset$. This is a contradiction to the assumption which implied $t' \in (R\downarrow_{[i-1]} \times \xi_{j^*}) \cap R$. □

5 Weak Orderings and (Half-)closed Intervals

So far, we did only consider root functions θ_1, θ_2 that are either strictly ordered or equal on some underlying cell in Theorems 2 and 5. We will now consider root functions that are ordered less-or-equal.

To motivate this, assume we want to construct a cell where some constraint $p \leq 0$ holds. Theorem 3 tells us that we can construct a cell around some sample s such that $p(s) \leq 0$ where p is sign-invariant by requiring that every root of p remains outside the cell that we describe. However, this yields a smaller cell guaranteeing the stronger property $p < 0$. Similarly, as in the previous section, we could merge this cell with the adjacent cells where $p = 0$ holds (again, as long as we cut these cells such that they are in the same cylinder).

This merger might not only produce bigger cells, it also avoids lifting over roots of polynomials, which is particularly expensive as non-rational number computations cannot be avoided in this case.

To describe such cells, we extend the notion of *symbolic intervals* to closed and half-closed intervals. For constructing such a cell, we will (1) generalize Theorems 3, 4 and 7 for half-closed intervals and weak orderings, and (2) define how we determine a symbolic interval that describes the cell as general as possible.

Now, we extend Theorem 3 allowing to realize semi-sign-invariance:

Theorem 6 (Root Ordering for Semi Sign Invariance). *Let* $i \in \mathbb{N}$, $R \subseteq \mathbb{R}^i$ *be connected,* $p \in \mathbb{Q}[x_1, \dots, x_{i+1}]$ *be irreducible such that* $\mathrm{level}(p) = i+1$, *and* $\mathtt{I}$ *be a symbolic interval of level* $i+1$. *Assume that* p, $\mathtt{I}.l.p$, *and* $\mathtt{I}.u.p$ *(if defined) are analytically delineable on* R.

Assume that for each real root function θ *of* p *on* R *it holds*

- *if* $\mathtt{I} = (\mathtt{I}.l, \mathtt{I}.u)$, *then either* $\mathtt{I}.l \neq -\infty$ *and* $\theta \leq \mathtt{I}.l$ *on* R, *or* $\mathtt{I}.u \neq \infty$ *and* $\mathtt{I}.u \leq \theta$ *on* R;
- *if* $\mathtt{I} = [\mathtt{I}.l, \mathtt{I}.u]$, *then either* $\mathtt{I}.l \neq -\infty$ *and* $\theta < \mathtt{I}.l$ *on* R, *or* $\mathtt{I}.u \neq \infty$ *and* $\mathtt{I}.u < \theta$ *on* R.

Then either $p < 0$ *on* R, *or* $p > 0$ *on* R.

Assume that for each real root function θ *of* p *on* R *it holds*

- *if* $\mathtt{I} = [\mathtt{I}.l, \mathtt{I}.u]$, *then either* $\mathtt{I}.l \neq -\infty$ *and* $\theta \leq \mathtt{I}.l$ *on* R, *or* $\mathtt{I}.u \neq \infty$ *and* $\mathtt{I}.u \leq \theta$ *on* R.

Then either $p \leq 0$ *on* R, *or* $p \geq 0$ *on* $R \times \mathtt{I}$. △

So far, this only exploits the sign conditions on the input polynomials. To carry this down to lower dimensions, we generalize and extend Theorem 5 for weak orderings:

Theorem 7 (Root Orderings for Pairs of Root Functions / Lifting of Pairs of Root Functions (Weak Case)). *Let* $i \in \mathbb{N}$, $R \subseteq \mathbb{R}^i$ *be a connected analytic submanifold,* $s \in R$, $p_1, p_2 \in \mathbb{Q}[x_1, \dots, x_{i+1}]$ *be irreducible and coprime such that* $\mathrm{level}(p_1) = \mathrm{level}(p_2) = i+1$, $\theta_1, \theta_2 : R \to \mathbb{R}$ *be real root functions of* p_1 *and* p_2 *respectively such that* $\theta_1(s) < \theta_2(s)$. *Let* $\mathtt{I}$ *be a symbolic interval of level* i, *and* $\mathrm{irExpr}(\mathrm{factors}(\mathrm{res}_{x_{i+1}}(p_1, p_2)), s_{[i-1]}) = \{\xi_1, \dots, \xi_k\}$ *s.t.* $\xi_1(s_{[i-1]}) \leq \dots \leq \xi_k(s_{[i-1]})$.

Assume that p_1 *and* p_2 *are analytically delineable on* R, $\mathtt{I}.l.p$ *and* $\mathtt{I}.u.p$ *are analytically delineable on* $R\downarrow_{[i-1]}$, $\{\mathrm{res}_{x_{i+1}}(p_1, p_2), \mathit{proj\_polys}(p_1), \mathit{proj\_polys}(p_2)\}$ *is locally* $(R\downarrow_{[i-1]} \times \mathtt{I})$*-delineable, and for every* $j \in [1..k]$, *if*

- $s' := \xi_j(s_{[i-1]}) \notin \mathit{del\_int}(\{p_1, p_2\}, (s_{[i-1]}, s'))$, *or*
- $s' := \xi_j(s_{[i-1]}) \in \mathit{del\_int}(\{p_1, p_2\}, (s_{[i-1]}, s'))$ *and* $\theta_1(s_{[i-1]}, s') = \theta_2(s_{[i-1]}, s')$,

then for each $\sim \in \{<, \leq\}$ *it holds*

- *if* $\mathtt{I} = (\mathtt{I}.l, \mathtt{I}.u)$, *then* $\mathtt{I}.l \neq -\infty$ *and* $\xi_j \leq \mathtt{I}.l$ *on* $R\downarrow_{[i-1]}$ *or* $\mathtt{I}.u \neq \infty$ *and* $\mathtt{I}.u \leq \xi_j$ *on* $R\downarrow_{[i-1]}$;

– *if* $\mathtt{I} = [\mathtt{I}.l, \mathtt{I}.u]$*, then* $\mathtt{I}.l \neq -\infty$ *and* $\xi_j \sim \mathtt{I}.l$ *on* $R\downarrow_{[i-1]}$ *or* $\mathtt{I}.u \neq \infty$ *and* $\mathtt{I}.u \sim \xi_j$ *on* $R\downarrow_{[i-1]}$.

Then $\theta_1 \sim \theta_2$ *holds on* R. △

The generalization to half-open intervals is straight-forward. We note that Theorem 4 can also be generalized for closed and half-closed intervals as well; as this is mainly technical, we omit it here.

We did not examine yet how to choose the symbolic interval. As this is mainly technical, we give a brief sketch of the algorithm: Throughout the algorithm, we maintain sets of polynomials that need to be made (semi-)sign-invariant or order-invariant and orderings of root functions respectively. These properties define *critical roots* at which their satisfaction might change, which are classified as *relevant* (root must not be inside the interval), *weakly relevant* (root might intersect with a closed bound of the interval) or *irrelevant* (root might be inside the cell). After we computed and classified these roots according to Theorems 4, 6 and 7, we choose an admissible symbolic interval (e.g. the largest one). Afterwards, we compute an ordering on the critical roots w.r.t. the chosen symbolic interval, and compute further projection polynomials e.g. for (local) delineability.

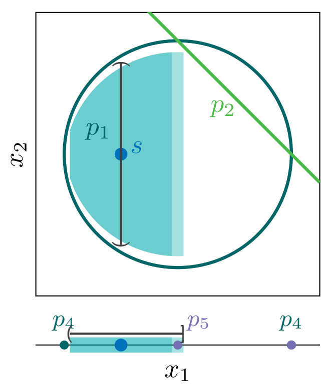

Fig. 6: The cell where two polynomials are sign-invariant can be made right-closed. The darker cell depicts the cell in traditional single cell construction, the lighter cell the bigger cell if we allow closed interval bounds.

Example 4. Consider the polynomials $p_1 = x_1^2 + x_2^2 - 1$ and $p_2 = -1 + x_1 + x_2$, and assume we construct the sign-invariant cell around $s = (-0.5, 0)$. Figure 6 depicts this example where a cell is constructed which is sign-invariant for p_1 and p_2. As the interval for x_2 has open bounds defined by p_1, the intersection point of p_1 and p_2 will be outside that interval, even if we extend the underlying cell in that direction. Thus, the root of the resultant p_5 is weakly relevant, and we can close the upper bound on x_1. We note that we cannot close the other bound for the reason that the indexed roots of the defining polynomials for the interval of x_2 need to be well-defined on the underlying cell. △

6 Experimental Results

We implemented our single cell construction algorithm in our solver `SMT-RAT`, which uses it for generating explanations for the NLSAT algorithm; as this implementation is incomplete due to McCallum's projection operator, we use the method as described in [6] as fallback. For evaluating the practical usability of the above results, we implemented several variants which vary in when and how to check for irrelevant roots. For the evaluation, we focus on the following variants:

Baseline A variant without any checks for irrelevant roots.

All For each root of a resultant that is inside the delineable interval of the two originating polynomials, we isolate their roots to determine whether there is an irrelevant intersection as specified by the root ordering to be maintained.

Independent As All, but we only check the roots of resultants which are irreducible. Further, we only make use of local delineability if all of its roots are not relevant; otherwise, we just require its order-invariance. These are exactly the cases analogous to Example 3 where we do not need to compute iterative resultants with the given resultant.

Rational As All, but we only check whether a root is irrelevant if it is rational and the underlying sample point does not contain a non-rational number. We thus focus on easy-to-compute checks, as root isolation of polynomials above non-rational samples is computationally expensive.

Bounds We do not check for irrelevant roots, but check whether a bound of a symbolic interval can be closed using weak root orderings as described in Section 5. We restrict to cases where no two root functions intersect over the current sample, as in those cases, there will be a section on the level below. Applying our optimization would require adding more discriminants to the projection, which is otherwise avoided by using theorems on equational constraints as described in [10].

There are more variants thinkable, e.g. deciding whether to check for irrelevance based on polynomial degrees or level. We justify in the evaluation below why we excluded those variants from the evaluation.

We conduct our experiments on Intel®Xeon®Platinum 8160 CPUs with 2.1GHz per core on SMT-LIB's *QF_NRA* benchmark set [11], using a time limit of 60 seconds and a memory limit of 4 Gigabyte. The code, instructions for reproducing and raw results are available at `https://doi.org/10.5281/zenodo.11071146`.

We restrict the following evaluation to the instances where the single cell construction algorithm is called at least once, leaving 7615 instances; the other instances are solved by Boolean reasoning. The results are depicted in Figure 7a and Figure 7b, clearly showing that All and Independent perform worse than

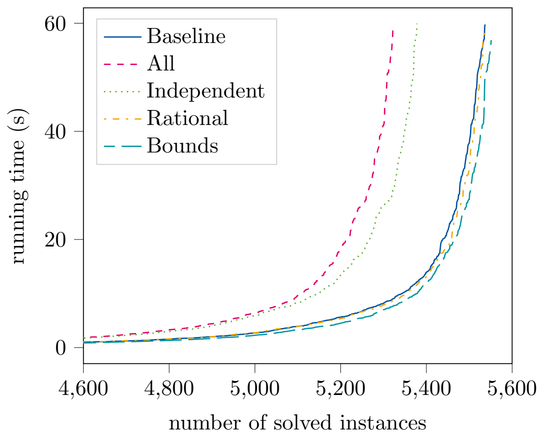

(a) Performance profile of all tested variants.

	gained	lost	diff
All	53	267	-214
Independent	9	167	-158
Rational	42	41	1
Bounds	85	70	15
virtual best	113	0	113

(b) Number of gained and lost instances and the total difference compared to Baseline. The virtual best is the portfolio of all variants.

Fig. 7: Results of all variants.

Baseline, while Rational and Bounds solve about as many as Baseline. Considering the number of gained and lost instances, All and Independent are truly inferior to Baseline; for Rational and Bounds, the numbers show typical differences for slightly different but similar configurations in SMT solving for NRA. In the following, we will elaborate why the first two variants are inferior, and how the behaviour of the latter two differs from Baseline. For reasons of comparability, this evaluation is based on the 5209 instances that are solved by all variants.

All is able to use about 3% of all resultant roots to merge adjacent cells. Merging happens in 1002 of the 5210 instances. On these instances, the quality of the generated cells *slightly* increases, which we measure by the time spent in the NLSAT engine, as shown in Table 1a. However, determining which roots are irrelevant is expensive: it takes about 25% of the overall running time.

Independent is able to use about only 0.08% of all resultant roots for merging (distributed over 60 instances); this is due to the low portion of resultants that are detected as independent and have at least one root (0.66% of all re-

Table 1: Influence of merges on the performance.

(a) All: Average time (in seconds) spent in the NLSAT engine (total running time minus time spent in the theory backends), split by instances where at least one merge happens and where no merges happen.

	> 0 merges	no merge
Baseline	1.08	0.21
All	0.9	0.21
# instances	1002	4208

(b) Bounds: Average number of sections (on any level), split by instances where at least one merge happens and where no merges happen.

	> 0 merges	no merge
Baseline	0.029	0.001
Bounds	0.023	0.001
# instances	4071	1139

sultants). The time spent on checking for irrelevant roots is as high as for All (about 25% of the total running time) - this might be due to the high portion of irreducible resultants (81%). Still, Independent loses fewer instances than All compared to Baseline: For explaining this, we look at the instances on which Independent and All are exceeding the time limit; for the former, this happens during checking for irrelevant roots on 344 instances, for the latter, it happens on 452 instances. There are two possible reasons: First, Independent only checks irreducible resultants. Second, Independent requires fewer resultants to be locally delineable, which might require more complex root orderings on the levels below. Tracing back the exact reasons however is not obvious.

Rational is able to use about only 0.7% of all resultant roots for merging (distributed over 351 instances); 32% of all roots are not considered as they or their underlying sample are non-rational. The time spent on checking for irrelevant roots is reduced to a negligible portion of the running time. Thus, the difference in behaviour compared to Baseline is only marginal.

Bounds is able to close one or both bounds of an interval in about 2% of the cases (distributed among 4071 instances). The above measure of quality for the generated cells (time spend outside the theory solver) changes only slightly. However, the number of sections decreases in instances where at least one merge happens, see Table 1b. A reason for the little effect might be that the interval bounds are more likely to be closed on higher levels: We observe that 12% of the intervals on the highest level are closed, 2.3% on the second highest, and at most 0.3% on most of the levels below.

7 Conclusion

We introduced a framework that allows to detect “irrelevant” roots of resultants which witness adjacent cells that can be merged during single cell construction. During merging, we still maintain cylindricity of the resulting cell. For doing so, we introduced the notion of *local delineability* which we require for the resultants. Using this notion, we can generate potentially larger cells. In some cases, the computation of certain iterated resultants can completely be avoided. We further introduced a “lightweight” bound-relaxing variant that allows to close some bounds of the constructed cell. All variants potentially reduce the amount of real root isolation calls in the future solving process.

However, regarding running times, the experimental results of our implementation do not confirm that these savings realize in NLSAT: Although some measures hint that the quality of the constructed cells is slightly improved, the checks determining which roots are irrelevant are expensive. Restricting the checks to efficient or promising cases could reduce this issue; however, the approaches we examined (avoiding non-rational number computations, or restricting to cases where the number of projection polynomials is reduced) simply detect too few irrelevant roots to be effective. Similarly, the bound-relaxing variant is applicable in too few cases. Still, the approach holds potential. Firstly, the experimental results are limited to the SMT-LIB’s benchmark set, which might not reflect the diversity of all NRA problems. Secondly, we conjecture that bigger cells reflect the problem's structure better; we did not yet evaluate the impact on the quality of results in quantifier elimination due to the limited scope of this paper.

Future work could improve these issues: Firstly, we need efficient checks for irrelevant roots, for instance validated numerics approaches in the style of [12]. Furthermore, this paper did only examine irrelevant roots of resultants; it is unclear whether a similar approach for discriminants is possible. For improving the applicability of the bound-relaxing variant, the roots of discriminants and coefficient could further be investigated: The work in [2] follows a similar idea but uses different reasoning, the two approaches might be unified to be applicable in more cases.

References

1. Ábrahám, E., Davenport, J.H., England, M., and Kremer, G.: 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). DOI: 10.1016/j.jlamp.2020.100633
2. Bär, P., Nalbach, J., Ábrahám, E., and Brown, C.W.: Exploiting Strict Constraints in the Cylindrical Algebraic Covering. In: Satisfiability Modulo Theories (SMT 2023) (2023). https://ceur-ws.org/Vol-3429/paper13.pdf
3. Brown, C.W.: Improved Projection for Cylindrical Algebraic Decomposition. Journal of Symbolic Computation 32(5) (2001). DOI: 10.1006/jsco.2001.0463
4. Brown, C.W., and Košta, M.: Constructing a Single Cell in Cylindrical Algebraic Decomposition. Journal of Symbolic Computation 70 (2015). DOI: 10.1016/j.jsc.2014.09.024
5. Collins, G.E.: Quantifier Elimination for Real Closed Fields by Cylindrical Algebraic Decomposition. In: Automata Theory and Formal Languages (1975). DOI: 10.1007/3-540-07407-4_17
6. Jovanović, D., and de Moura, L.: Solving Non-linear Arithmetic. In: Automated Reasoning (IJCAR 2012) (2012). DOI: 10.1007/978-3-642-31365-3_27
7. Kremer, G., and Nalbach, J.: Cylindrical Algebraic Coverings for Quantifiers. In: Satisfiability Checking and Symbolic Computation 2022 (SC-square 2022) (2023). https://ceur-ws.org/Vol-3458/paper1.pdf
8. McCallum, S.: An Improved Projection Operation for Cylindrical Algebraic Decomposition. University of Wisconsin-Madison (1985).
9. McCallum, S.: An Improved Projection Operation for Cylindrical Algebraic Decomposition. In: Quantifier Elimination and Cylindrical Algebraic Decomposition (1998). DOI: 10.1007/978-3-7091-9459-1_12
10. Nalbach, J., Ábrahám, E., Specht, P., Brown, C.W., Davenport, J.H., and England, M.: Levelwise Construction of a Single Cylindrical Algebraic Cell. Journal of Symbolic Computation 123 (2024). DOI: 10.1016/j.jsc.2023.102288
11. Preiner, M., Schurr, H.-J., Barrett, C., Fontaine, P., Niemetz, A., and Tinelli, C.: *SMT-LIB Release 2023 (Non-Incremental Benchmarks)*, (2024). DOI: 10.5281/zenodo.10607722. 2024.
12. Strzeboński, A.: CAD Adjacency Computation Using Validated Numerics. In: International Symposium on Symbolic and Algebraic Computation (ISSAC '17) (2017). DOI: 10.1145/3087604.3087641