Computing the Dimension of Real Algebraic Sets

Pierre Lairez
Inria
France
pierre.lairez@inria.fr

Mohab Safey El Din
Sorbonne Université, CNRS, LIP6, Équipe PolSys
France
mohab.safey@lip6.fr

ABSTRACT

Let V be the set of real common solutions to $F = (f_1, \dots, f_s)$ in $\mathbb{R}[x_1, \dots, x_n]$ and D be the maximum total degree of the f_i's. We design an algorithm which on input F computes the dimension of V. Letting L be the evaluation complexity of F and $s = 1$, it runs using $O^{\sim}\left(LD^{n(d+3)+1}\right)$ arithmetic operations in $\mathbb{Q}$ and at most $D^{n(d+1)}$ isolations of real roots of polynomials of degree at most D^n.

Our algorithm depends on the *real* geometry of V; its practical behavior is more governed by the number of topology changes in the fibers of some well-chosen maps. Hence, the above worst-case bounds are rarely reached in practice, the factor D^{nd} being in general much lower on practical examples. We report on an implementation showing its ability to solve problems which were out of reach of the state-of-the-art implementations.

CCS CONCEPTS

• **Computing methodologies → Algebraic algorithms**.

KEYWORDS

Computer algebra; semi-algebraic set; dimension

ACM Reference Format:
Pierre Lairez and Mohab Safey El Din. 2021. Computing the Dimension of Real Algebraic Sets. In *Proceedings of the 2021 International Symposium on Symbolic and Algebraic Computation (ISSAC '21), July 18–23, 2021, Virtual Event, Russian Federation.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3452143.3465551

1 INTRODUCTION

Polynomial system solving over the reals finds numerous applications in engineering sciences. In mechanism design, mobility properties often translate to identifying the dimension of a real algebraic set, that is the real solution set to polynomial equations with real coefficients. This paper focuses on the design of an algorithm for computing dimension whose practical performance improves upon previously known algorithms.

Prior results. The cylindrical algebraic decomposition algorithm [11] computes a partition of the semi-algebraic set defined by an input semi-algebraic formula into finitely many cells which are homeomorphic to $]0,1[^i$ for some integer i. This integer i is the dimension of the corresponding cells. The dimension of the whole semi-algebraic set is the maximum dimension of its cells. The computational complexity of this approach is doubly exponential with respect to the dimension of the ambient space, denoted n.

Another approach [5; 24; 25] makes the most of the following characterization of the dimension: the dimension of a semi-algebraic set $S \subseteq \mathbb{R}^n$ is the maximum integer d such that the projection of S on some coordinate subspace of dimension d has non-empty interior. Quantifier elimination over the reals computes a semi-algebraic description of the projection of a semi-algebraic set. Hence, using quantifier elimination algorithms and their complexity analysis, one can derive the following complexity result. When the input is a single polynomial of total degree D in $\mathbb{R}[x_1, \dots, x_n]$, computing the dimension d of its real vanishing set is performed with $D^{O(d(n-d))}$ arithmetic operations in $\mathbb{R}$. Note that several polynomial constraints $f_1, \dots, f_s$ can be substituted by the single constraint $\sum_i f_i^2$, although in practice we try not to. A probabilistic variant refines this complexity bound to $O^{\sim}\left(n^{16}(D+1)^{3d(n-d)+5n+5}\right)$ arithmetic operations and enables practical implementations [4].

The dimension of a semi-algebraic set S is also the Krull dimension of its real radical [12; 31]. Such methods compute a prime decomposition of the radical ideal of the polynomials vanishing on S and decide whether associated irreducible components contain regular real points, yielding a smooth point in S. If not, one needs to investigate the singular loci of the algebraic sets defined by these components. Following this idea of computing smooth points in S, a dedicated numerical routine procedure for computing smooth points in semi-algebraic sets based on numerical homotopy and deflation algorithms is designed in [22]. Complexity bounds are worse than the state-of-the-art bounds but the numerical flavour of this algorithm makes it practically fast on well-conditioned examples.

Lastly, the related problem of computing the local dimension at a point has also been considered [32].

Problem statement. Current state-of-the-art algorithms are not satisfactory from a practical view: the best algorithms, from a complexity viewpoint, still rely on quantifier elimination which involves the computation of algebraic or semi-algebraic formulas whose sizes are $D^{O(d(n-d))}$. This proves to be a bottleneck. New numerical procedures have emerged but implementations are vulnerable to accuracy issues and hit the problem of deciding equality to zero from approximations.

The goal of this paper is to design a new algorithm for computing the dimension of a *real algebraic set*, practically efficient, which computes (and stores) algebraic data of size bounded by $(sD)^{O(n)}$.

Main results. Our algorithm takes as input polynomials $f_1, \dots, f_s$ in $\mathbb{R}[x_1, \dots, x_n]$. Let $D = \max_i \deg(f_i)$ and $V \subseteq \mathbb{R}^n$ be the real algebraic set defined by $f_1 = \cdots = f_s = 0$. For a sufficiently generic function $h \in \mathbb{R}[x_1, \dots, x_n]$, we show that we can compute finitely

ISSAC '21, July 18–23, 2021, Virtual Event, Russian Federation

ACM ISBN 978-1-4503-8382-0/21/07...$15.00
https://doi.org/10.1145/3452143.3465551

many points $t_0, \ldots, t_N \in \mathbb{Q}$ such that

$$\dim V = \max_{0 \le i \le N} \dim \left(V \cap h^{-1}(t_i) \right) + 1, \tag{1}$$

unless V is empty. When V is compact, we choose h to be a linear form. This leads to an induction on the number of variables. In the general case, we choose h to be the squared distance to a generic point. This reduces the general case to the compact case.

When $s = 1$, the points t_i are chosen so that $t_{i-1} < v_i < t_{i+1}$ where the points $v_i \in \mathbb{R}$ are limits of critical values of the restriction of h to the set $\{f = \varepsilon\}$ as $\varepsilon \to 0$. In particular, we prove that $N \le D^n$. The general case is more involved but analogue.

At each recursive call, the cost of the algorithm lies in:

(i) deciding $V \neq \emptyset$, in order to apply (1);
(ii) computing a univariate polynomial whose zero set contains $\{v_1, \ldots, v_N\}$;
(iii) isolating the zeros of this polynomial to find appropriate points t_i.

The steps (i) and (ii) are performed in $(sD)^{O(n)}$ arithmetic operations. The worst case complexity of the last step depends on the height of the coefficients of the input equations; we will not enter into such considerations. The total number of recursive calls is $(sD)^{O(nd)}$.

Our results do not improve on the state-of-the-art asymptotic complexity bound which is exponential in $d(n-d)$ while we only reach nd (and not counting the cost of isolating the real roots of some univariate polynomials). However, the behavior of the algorithm depends more on the *real geometry* of V than what is observed for algorithms based on quantifier elimination. Nonetheless, the excellent practical behavior is easily explained.

For one part, the practical behavior of the algorithm is governed by the actual degree of the algebraic varieties arising in the computation of the limits of critical values. It is typically lower than the worst-case bounds. For another part, the practical behavior is governed by the number of recursive calls, which is $(sD)^{O(nd)}$ where d is the dimension of V. With previous notations, if V is not empty, there are $N+1$ direct recursive calls (and each one may induce other indirect recursive calls naturally). While the polynomial computed in step (ii) has degree at most $(6sD)^n$, the actual degree is often lower than this bound. And the number N of *real* roots is even lower (see §4). In our examples, the total number of recursive calls ends up being dramatically lower than the worst-case bound. This accounts for the good practical performance of our algorithm.

We report in detail on practical experiments which illustrate the interest of our approach. Our algorithm is able to solve problems which are out of reach of the state-of-the-art implementations. We give detailed information on its execution, in particular the maximum number N observed at each level of the recursion.

Related works. The study of properties of critical loci of generic quadratic or linear maps that can be exploited for polynomial system solving is initiated in [1; 2] and followed by [3]. Properness properties are already exploited in algorithms of real algebraic geometry in [30]. The complexity analysis uses results about the complexity of the geometric resolution algorithm [19, and references therein]. The design of computer algebra algorithms whose practical behavior depends more on the real geometry of the set under study is already exploited in [28] for answering connectivity queries (see also [9] for more recent developments).

2 GEOMETRY

Let $F = (f_1, \ldots, f_s)$ be a polynomial map $\mathbb{R}^n \to \mathbb{R}^s$. In this whole section, we denote by $V \subset \mathbb{R}^n$ the real algebraic set $F^{-1}(0)$. Let $h \in \mathbb{R}[x_1, \ldots, x_n]$ be another polynomial. For $t \in \mathbb{R}$, let $V(t) = V \cap h^{-1}(t)$. We study the dimension of V in terms of the dimension of the fiber $V(t)$.

2.1 Dimension of fibers

The dimension of V can be related to the dimension of the fibers $V(t)$. In the following statements, we use the convention $\dim \emptyset = -1$.

Proposition 2.1. *Let $Z \subset \mathbb{R}$ be a finite set such that $t \mapsto \dim V(t)$ is locally constant on $\mathbb{R} \setminus Z$. If $V \neq \emptyset$, then*

$$\dim V = \max \left(\max_{r \in Z} \dim V(r), \max_{t \in \mathbb{R} \setminus Z} \dim V(t) + 1 \right).$$

Proof. By Hardt's triviality theorem [21], there is a finite set $Z' \subset \mathbb{R}$ such that h induces a semialgebraic locally trivial fibration on $\mathbb{R} \setminus Z'$: if U is a connected component of $\mathbb{R} \setminus Z'$, then $V \cap h^{-1}(U)$ is isomorphic, as a semialgebraic set, to $U \times V(t)$, for any $t \in U$. Since V is the finite union of all $V(r)$, for $r \in Z'$, and all $V \cap h^{-1}(U)$, for all connected components U of $\mathbb{R} \setminus Z'$, we obtain that

$$\dim V = \max \left(\max_{r \in Z'} \delta(r), \max_{U} \dim(V \cap h^{-1}(U)) \right).$$

Since $V \cap h^{-1}(U) \simeq U \times V(t)$ for any $t \in U$, we have $\dim(V \cap h^{-1}(U)) = \delta(t) + 1$, unless it is empty. When it is empty, then $\delta(t) + 1 = 0$ holds, and since $\dim V \ge 0$, it holds that

$$\dim V = \max \left(\max_{r \in Z'} \delta(r), \max_{t \in \mathbb{R} \setminus Z'} \delta(t) + 1 \right), \tag{2}$$

which is the claim, albeit with Z' in place of Z.

Without loss of generality, we may assume that $Z \subseteq Z'$. We now show that we can remove points from Z' without breaking Equation (2) as long as $\dim V(t)$ is locally constant on $\mathbb{R} \setminus Z$. If $Z = Z'$, there is naturally nothing to prove. Assume that $Z' = Z \cup \{a\}$ for some $a \notin Z$ (and the general case follows by induction). Let $\delta(t)$ denote $\dim V(t)$. Since δ is locally constant on $\mathbb{R} \setminus Z$, and $a \notin Z$, δ is constant in a neighborhood of a. So

$$\max_{t \in \mathbb{R} \setminus Z'} \delta(t) + 1 = \max_{t \in \mathbb{R} \setminus Z} \delta(t) + 1.$$

Combining with (2), this gives

$$\dim V = \max \left(\max_{r \in Z} \delta(r), \delta(a), \max_{t \in \mathbb{R} \setminus Z} \delta(t) + 1 \right).$$

To conclude, we observe that

$$\delta(a) \le \delta(a) + 1 \le \max_{t \in \mathbb{R} \setminus Z} \delta(t) + 1$$

because $a \in \mathbb{R} \setminus Z$. □

2.2 Generic case

Under a genericity assumption on h, the special fibers above Z in Proposition 2.1 carry no information.

Let $\mathcal{L}$ be the set of linear forms on $\mathbb{R}^n$ and let Q be the set of quadratic functions $x \mapsto \|x - p\|^2$ for $p \in \mathbb{R}^n$. Both are real algebraic varieties isomorphic to $\mathbb{R}^n$.

Proposition 2.2. *Let $S \subset \mathbb{R}^n$ be a semialgebraic set of positive dimension. If h is a generic element of $\mathcal{L}$ or $\mathcal{Q}$, then for any $t \in \mathbb{R}$,*

$$\dim\left(S \cap h^{-1}(t)\right) < \dim(S).$$

Proof. Let $d = \dim S$. Let W be the Zariski closure of S in $\mathbb{R}^n$. Note that $\dim W = \dim S$ [7, Proposition 2.8.2]. It is enough to prove that for any irreducible component Y of W,

$$\dim\left(Y \cap h^{-1}(t)\right) < d. \tag{3}$$

Indeed, by the inclusion $S \subseteq W$, and the formula for the dimension of a union [7, Proposition 2.8.5(i)], we have

$$\dim\left(S \cap h^{-1}(t)\right) \leq \max_Y \dim\left(Y \cap h^{-1}(t)\right).$$

Let Y be an irreducible component of W. We may assume that $\dim Y \geq 1$ as (3) is trivial otherwise (because $d \geq 1$). In particular, Y contains at least two points p and q. When $h \in \mathcal{L}$ (resp. $\mathcal{Q}$) is generic, it is clear that $h(p) \neq h(q)$. In particular, h is not constant on Y and therefore $h - t$ is a nonzero function on Y for any $t \in \mathbb{R}$. By irreducibility of Y and Krull's principal ideal theorem, it follows that $\dim\left(Y \cap h^{-1}(t)\right) < \dim Y$, which gives (3) and concludes the proof. □

Proposition 2.3. *Assume that h is a generic element in $\mathcal{L}$ or $\mathcal{Q}$. Let $Z \subset \mathbb{R}$ be a finite set such that $t \mapsto \dim V(t)$ is locally constant on $\mathbb{R} \setminus Z$. If $V \neq \emptyset$, then*

$$\dim V = \max_{t \in \mathbb{R}\setminus Z} \dim V(t) + 1.$$

Proof. Proposition 2.2 implies that $\dim V > \max_{r \in Z} \dim V(r)$. We conclude with Proposition 2.1. □

2.3 Limits of critical values

The main obstacle in applying Propositions 2.1 or 2.3 is the computation of a suitable finite set $Z \subset \mathbb{R}$ such that the dimension of the fiber $V(t)$ is locally constant outside Z. The proof of Proposition 2.1 indicates that an effective Hardt trivialization would be enough. Actually, much less is required. We show now how to compute such a finite set Z by means of critical points computations.

We start by introducing some notation. For $e \in \mathbb{R}^s$, let S_e be the semi-algebraic set

$$S_e = \left\{ p \in \mathbb{R}^n \;\middle|\; |f_i(p)| \leq e_i \text{ for } 1 \leq i \leq s \right\}.$$

For $t \in \mathbb{R}$, let $S_e(t) = S_e \cap h^{-1}(t)$.

We will see that the dimension of $V(t)$ is entirely determined by the homotopy type of $S_e(t)$, for a generic and small enough $e \in \mathbb{R}^s$. Moreover, the variations of the homotopy types of $S_e(t)$ are controlled by Thom's isotopy lemma when t varies in terms of the critical values of h on a Whitney stratification of $S_e(t)$. Such a stratification is easy to get when e is generic.

In all this section, we assume that for e in a neighborhood of 0, the restriction of h to S_e is proper. In particular, this is the case when F is proper on $\mathbb{R}^n$ (in which case S_e is compact), or when h is proper on $\mathbb{R}^n$.

Let $e \in \mathbb{R}^s$ and let H_e be the hypercube $[-e_1, e_1] \times \cdots \times [-e_s, e_s]$ (if $e_i < 0$, then $[-e_i, e_i]$ means $[e_i, -e_i]$). The relative interior of the 3^s facets of H_e form a Whitney stratification, which we denote $\mathcal{WH}_e$; namely

$$\mathcal{WH}_e = \left\{ I_1 \times \cdots \times I_s \;\middle|\; I_i \in \left\{ \{-e_i\},]-e_i, e_i[, \{e_i\} \right\} \right\}.$$

When e is generic, the pullback $\mathcal{WS}_e = F^{-1}(\mathcal{WH}_e)$ induces a Whitney stratification of S_e. For this, it is enough to check that F is *transverse* to each strata of $\mathcal{WH}_e$ [18, (1.4)]. The map F is transverse to a submanifold $A \subseteq \mathbb{R}^s$ if for any $x \in F^{-1}(A)$,

$$\mathrm{d}_x F(\mathbb{R}^n) + T_{f(x)} A = \mathbb{R}^s. \tag{4}$$

Let $\mathcal{E} \subseteq \mathbb{R}^s$ be the set of all e such that $\mathcal{WS}_e$ is a Whitney stratification of S_e. The set $\mathcal{E}$ contains a dense Zariski open set in $\mathbb{R}^s$ [8, General position lemma]. In particular, if $e \in \mathcal{E}$, then $se \in \mathcal{E}$ for all but finitely many $s \in \mathbb{R}$.

For $e \in \mathbb{R}^s$, let Σ_e be

$$\Sigma_e = \cup_{A \in \mathcal{WS}_e} h\big(\operatorname{crit}(h, A^{\text{reg}})\big),$$

where A^{reg} denotes the regular locus of A (the points of A where the derivatives of the equations defining A are linearly independent). When $\mathcal{WS}_e$ is a Whitney stratification, then, by definition, the strata as regular and $A = A^{\text{reg}}$, so that

$$\Sigma_e = \cup_{A \in \mathcal{WS}_e} h\big(\operatorname{crit}(h, A)\big).$$

By Sard's theorem, Σ_e is finite. We may consider the limit $\lim_{r \to 0} \Sigma_{re}$. It is the finite set

$$\lim_{r \to 0} \Sigma_{re} = \cap_{R > 0} \overline{\cup_{0 < r < R} \Sigma_{re}}, \tag{5}$$

where the bar denotes closure with respect to Euclidean topology. This core construction will be heavily used in the algorithms we describe in the next section.

Lemma 2.4. *Two real compact algebraic sets that are homotopically equivalent have the same dimension.*

Proof. Since the homology is invariant under homotopy equivalence [23, Corollary 2.11], it is enough to give a homological characterization of dimension.

Let W be a compact real algebraic set. We claim that the dimension of W is the largest integer d such that $H_d(W; \mathbb{Z}_2)$ is not zero (where $\mathbb{Z}_2 = \mathbb{Z}/2\mathbb{Z}$). Indeed, W is homeomorphic to a d-dimensional simplicial complex [7, Theorem 9.2.1]. In particular, $H_k(W; \mathbb{Z}_2) = 0$ for $k > d$. Moreover, the sum of the d-dimensional simplices of a given simplicial decomposition of W is a nonzero element in $H_d(W; \mathbb{Z}_2)$ [7, Proposition 11.3.1], and then $H_d(W; \mathbb{Z}_2) \neq 0$. □

The next result establishes that we can take $Z = \lim_{r \to 0} \Sigma_{re}$.

Theorem 2.5. *Let $e \in \mathcal{E}$ such that the restriction of h to S_e is proper. The dimension of $V(t)$ is locally constant on $\mathbb{R} \setminus \lim_{r \to 0} \Sigma_{re}$.*

Proof. Let $U \subseteq \mathbb{R}$ be a closed interval such that $U \cap \lim_{r \to 0} \Sigma_{re}$ is empty. There exists $0 < r < 1$ such that $re \in \mathcal{E}$ and $\Sigma_{re} \cap U = \emptyset$. For notations, we may replace e by re and assume that $r = 1$.

The sets $A \cap h^{-1}(U)$, for $A \in \mathcal{WS}_e$, form a Whitney stratification of $S_e \cap h^{-1}(U)$. By construction, $\operatorname{crit}(h, A \cap h^{-1}(U))$ is empty. Indeed $e \in \mathcal{E}$, so $\mathcal{WS}_e$ is a Whitney stratification and $A = A^{\text{reg}}$. Moreover,

$$\begin{aligned} \operatorname{crit}\left(h, A^{\text{reg}} \cap h^{-1}(U)\right) &= \operatorname{crit}(h, A^{\text{reg}}) \cap h^{-1}(U) \\ &\subseteq h^{-1}\big(\Sigma_e \cap U\big) = \emptyset. \end{aligned}$$

So the restriction of h to each $A \cap h^{-1}(U)$ is a submersion. Besides, the restriction of h to S_e is proper, by hypothesis. Therefore, by Thom's first isotopy lemma [27, Proposition 11.1], h induces a locally trivial fibration over U. Since U is an interval, it is simply connected

Algorithm 1 Checking the genericity of the stratification

Input $f_1, \ldots, f_s \in \mathbb{R}[x_1, \ldots, x_n]$, $e \in \mathbb{R}^s$
Output TRUE or FALSE
Postcondition If TRUE is returned, then $\mathcal{WS}_e$ is a Whitney stratification of S_e. If e is generic, then returns TRUE.

```
function CheckWhitneyStratification(f_1, ..., f_s, e)
    for I ⊆ {1, ..., s} with #I ≤ n + 1 and σ : I → {±1} do
        M ← (∂f_i/∂x_j)_{i∈I, j∈{1,...,n}}, a #I × n matrix
        J ← ⟨f_i − σ_i e_i, i ∈ I⟩ + ⟨#I × #I minors of M⟩
        if J ≠ ⟨1⟩ then
            return FALSE
    return TRUE
```

and the fibration is trivial. If we choose a base point $b \in U$, this means that there is a diffeomorphism $\phi : S_e \cap h^{-1}(U) \to U \times S_e(b)$ such that $\mathrm{proj}_1 \circ \phi = h$. It follows easily that all $S_e(t)$, for $t \in U$ are diffeomorphic and, in particular, homotopically equivalent.

Next, for any $t \in U$, the inclusion $V(t) \to S_e(t)$ is a homotopy equivalence [13, Proposition 1.6]. So for any $t, s \in U$, we have the homotopy equivalence

$$V(t) \underset{\text{Durfee}}{\simeq} S_e(t) \underset{\text{Thom}}{\simeq} S_e(s) \underset{\text{Durfee}}{\simeq} V(s).$$

So the homotopy type of $V(t)$ is constant for $t \in U$. By Lemma 2.4, the dimension of $V(t)$ is constant for $t \in U$. □

3 ALGORITHMS

An algorithm for computing the dimension mostly follows from Proposition 2.3 and Theorem 2.5. We first explain how to check a sufficient condition for e to be in $\mathcal{E}$.

3.1 Genericity of the stratification

We use the notations of §2.3. To check that $\mathcal{WS}_e$ is a Whitney stratification of S_e, it is enough to check Condition (4) for any strata A of $\mathcal{WH}_e$. In order to reduce to easier algebraic computations, we use instead the Zariski closure of the strata and check the condition over $\mathbb{C}$. This gives a stronger set of conditions, that are are still sufficient conditions for $\mathcal{WS}_e$ to be a Whitney stratification.

A stratum A of $\mathcal{WH}_e$ is determined by a subset $I \subseteq \{1, \ldots, s\}$ and a map $\sigma : I \to \{\pm 1\}$. I and σ define a linear variety $L_{I,\sigma} = \{y \in \mathbb{C}^s \mid y_i = \sigma_i e_i, i \in I\}$ which is the Zariski closure of the corresponding strata of $\mathcal{WH}_e$, namely $L_{I,\sigma} \cap H_e$. The tangent space of $L_{I,\sigma}$ at any $y \in L_{I,\sigma}$ is $\{y \in \mathbb{C}^s \mid y_i = 0, i \in I\}$. Therefore, if p_I is the linear projection $\mathbb{R}^n \to \mathbb{R}^I$ retaining only coefficients whose indices are in I, Condition (4) is equivalent to $p_I \circ \mathrm{d}_x F : \mathbb{R}^n \to \mathbb{R}^I$ being surjective. Therefore, if the variety

$$Y_{I,\sigma} = \left\{ x \in \mathbb{C}^n \;\middle|\; F(x) \in L_{I,\sigma} \text{ and } p_I \circ \mathrm{d}_x F \text{ not surjective} \right\}.$$

is empty, then Condition (4) holds for the strata defined by I and σ. The equations defining $Y_{I,\sigma}$ are $f_i = \sigma_i e_i$, for $i \in I$, and the vanishing of the $\#I \times \#I$ minors of $p_I \circ \mathrm{d}_x F$ (which is identified with a $\#I \times n$ matrix with polynomial coefficients).

There are 3^s possible values for I and σ. However, when $s > n+1$, not all $Y_{I,\sigma}$ need to be checked for emptyness. Indeed, when $\#I > n$, the map $p_I \circ \mathrm{d}_x F$ is never surjective (and the corresponding matrix has no $\#I \times \#I$ minors). In particular, $Y_{I,\sigma}$ contains $Y_{I',\sigma'}$ for any I'

Algorithm 2 Limits of critical values

Input $f_1, \ldots, f_s$ and $h \in \mathbb{R}[x_1, \ldots, x_n]$ and $e \in (\mathbb{R} \setminus \{0\})^s$
Precondition The restriction of h to $S_e = \{|f_i| \leq |e_i|\}$ is proper.
Output A finite set $Z \subset \mathbb{R}$ description as the zero set of a univariate polynomial
Postcondition $\lim_{r\to 0} \cup_{A \in \mathcal{WS}_{re}} h\left(\mathrm{crit}(h, A^{\mathrm{reg}})\right) \subseteq Z$

```
function LimitCriticalValues(f_1, ..., f_s, h, e)
    Z ← ∅
    for I ⊆ {1, ..., s} with #I ≤ n and σ : I → {±1} do
        J' ← ⟨Σ_{i∈I} λ_i ∂_j f_i − ∂_j h⟩_{j∈[n]} + ⟨σ_j e_j f_i − σ_i e_i f_j⟩_{i,j∈I}
            ▷ The λ_i are new variables. This is the ideal of W'_{I,σ,e}.
        J ← J' ∩ R[x_1, ..., x_n] + ⟨f_1, ..., f_s⟩
                                            ▷ Ideal of V ∩ closure(W_{I,σ,e})
        p ← a generator of (J + ⟨h − t⟩) ∩ R[t]
        Z ← Z ∪ p^{-1}(0)
    return Z
```

and σ' such that $I \subseteq I'$ and $\sigma = \sigma'|_I$. Therefore, to check that $Y_{I,\sigma} = \varnothing$ for all I and σ, it is enough to check it for all I with $\#I \leq n+1$. This leads to Algorithm 1 and the following statement.

Proposition 3.1. *On input $f_1, \ldots, f_s \in \mathbb{R}[x_1, \ldots, x_n]$ and $e \in \mathbb{R}^s$, Algorithm 1 returns* TRUE *if e is generic. If it returns* TRUE*, then $\mathcal{WS}_e$ is a Whitney stratification of S_e.*

3.2 Computing the limits of critical values

3.2.1 Description. We consider the problem of computing the limit set $\lim_{r\to 0} \Sigma_{re}$, for a given $e \in (\mathbb{R} \setminus \{0\})^s$, using the notations of §2.3. To this purpose, we extend to the case of several equations the method of [29], and remove the genericity assumptions. Most of the technical difficulties come from considering several equations.

Let P_e denote the critical points of h on the strata of $\mathcal{WS}_e$, that is $P_e = \cup_{A \in \mathcal{WS}_e} \mathrm{crit}(h, A^{\mathrm{reg}})$, so that $\Sigma_e = h(P_e)$. The limit set $\lim_{r\to 0} P_{re}$ is defined as in (5).

Lemma 3.2. *Let $e \in \mathbb{R}^s$ such that h is proper on S_e. Then*

$$\lim_{r\to 0} \Sigma_{re} = h\big(\lim_{r\to 0} P_{re}\big).$$

Proof. This commutation of h and $\lim_{r\to 0}$ is an elementary property of proper maps. The right-to-left inclusion follows directly from the definitions. Conversely, let $y \in \lim_{r\to 0} \Sigma_{re}$. Let $(R_i)_{i\geq 1}$ be a sequence decreasing to 0, with $R_1 \leq 1$. For any $i \geq 1$, $y \in \overline{\cup_{0<r<R_i} h(P_{re})}$, by definition. So there is some $x_i \in \cup_{0<r<R_i} P_{re}$ such that $|y - h(x_i)| \leq 1/i$. Note that $P_{re} \subseteq S_{re}$, so $x_i \in S_e$ (using $R_i \leq 1$). In particular, $x_i \in h^{-1}([y-1, y+1]) \cap S_e$, which is a compact set, by hypothesis. Up to extracting a subsequence, we may assume that $(x_i)_i$ converges to some $x \in \mathbb{R}^n$. By continuity, $h(x) = y$. Since $R_j \leq R_i$ for $j \geq i$, it follows that $x \in \overline{\cup_{0<r<R_i} P_{re}}$ for any $i \geq 1$. This means that $x \in \lim_{r\to 0} P_{re}$. □

Let Q_e be the complex analogue of P_e, that is the union over all I and σ of the complex critical points of h restricted to the regular locus of the complex stratum $F^{-1}(L_{I,\sigma})$. In other words,

$$Q_e = \cup_{I,\sigma} \,\mathrm{crit}\left(h, \{f_i = \sigma_i e_i \text{ for } i \in I\}^{\mathrm{reg}}\right) \subseteq \mathbb{C}^n.$$

It is clear that $P_e \subseteq Q_e$. Moreover, let $W_e = \cup_{r \in \mathbb{C}^*} Q_{re}$. If $x \in Q_{re}$, then r is entirely determined by x and e. Namely, if x lies in the stratum $F^{-1}(L_{I,\sigma})$, then $r = \sigma_i f_i(x)/e_i$, for any $i \in I$. This gives a well-defined regular map $\rho : W_e \to \mathbb{C}$ such that $x \in Q_{\rho(x)e}$ for any $x \in W_e$.

Proposition 3.3. *For any $e \in (\mathbb{R} \setminus \{0\})^s$, the set $h(V \cap \overline{W_e})$ is finite. Moreover, if h is proper on S_e, then $\lim_{r\to 0} \Sigma_{re} \subseteq h(V \cap \overline{W_e})$.*

Proof. We first prove the inclusion, when h is proper on S_e. Let $y \in \lim_{r\to 0} \Sigma_{re}$. By Lemma 3.2, there is some $x \in \lim_{r\to 0} P_{re}$ such that $y = h(x)$. By definition, $x \in \overline{\cup_{0<r<R} P_{re}}$ for any $R > 0$. Since $P_{re} \subseteq S_{re}$, it follows that $|f_i(x)| \leq R|e_i|$ for any $R > 0$ and any i. Therefore, $f_i(x) = 0$ for any i, and $x \in V$. Moreover, since $P_{re} \subseteq Q_{re} \subseteq W_e$, we have $x \in \overline{W_e}$.

For the finiteness, consider the map $\psi : x \in W_e \to \mathbb{C}^2$ defined by $\psi(x) = (\rho(x), h(x))$. Let $C = \psi(W_e)$. For any $r \in \mathbb{C}$, the set $C \cap \{z_1 = r\}$ is finite (where z_1 and z_2 are the coordinates on $\mathbb{C}^2$). Indeed,

$$C \cap \{z_1 = r\} = \psi\left(W_e \cap \rho^{-1}(r)\right) = \{r\} \times h\left(Q_{re}\right),$$

and Sard's theorem implies that $h(Q_{re})$ is finite. It follows that the set $\overline{C} \cap \{z_1 = r\}$ are finite too. (Otherwise $\overline{C}$ would contain a vertical line $\{z_1 = r\}$ and so C, which is dense in all the components of $\overline{C}$, would contain a dense subset of this line, which would contradict the finiteness of $C \cap \{z_1 = r\}$.)

Lastly, we observe that

$$\begin{aligned} h\big(V \cap \overline{W_e}\big) &= h\big(\overline{W_e} \cap \rho^{-1}(0)\big) = \operatorname{proj}_2\left(\psi\big(\overline{W_e}\big) \cap \{z_1 = 0\}\right) \\ &\subseteq \operatorname{proj}_2\left(\overline{C} \cap \{z_1 = 0\}\right), \end{aligned}$$

which is finite. □

Based on this statement, we give two ways of computing a finite set containing $\lim_{r\to 0} \Sigma_{re}$. One with Gröbner basis, without genericity hypothesis, and another using geometric resolution under genericity hypothesis, in order to obtain a complexity estimate.

3.2.2 Using Gröbner bases. For I and σ defining a stratum of $\mathcal{WS}_e$, consider the algebraic subvariety of $\mathbb{C}^n \times \mathbb{C}^{\#I}$

$$W'_{I,\sigma,e} = \big\{(x,\lambda) \mid \mathrm{d}_x h = \textstyle\sum_{i\in I} \lambda_i \mathrm{d}_x f_i \text{ and } \forall i,j \in I, \sigma_i e_i f_j = \sigma_j e_j f_i\big\}.$$

Let also $W'_e = \cup_{I,\sigma} W'_{I,\sigma,e}$.

Lemma 3.4. *For any $e \in (\mathbb{C} \setminus \{0\})^s$, $W_e = \operatorname{proj}_1(W'_e)$.*

Proof. The left-to-right inclusion follows directly from the definitions. Conversely, let $(x,\lambda) \in \cap W'_{I,\sigma,e}$. The equations $\sigma_i e_i f_j = \sigma_j e_j f_i$ imply the existence of a unique $r \in \mathbb{C}$ such that $f_i(x) = \sigma_i r e_i$ for all $i \in I$. Furthermore, we can choose λ such that the number of nonzero λ_i is minimized. Let $J = \{i \in I \mid \lambda_i \neq 0\}$. By minimality of J, the derivatives $\mathrm{d}_x f_i$, for $i \in J$, are linearly independent. In particular x is a regular point of the complex stratum $\{f_i = \sigma_i r e_i \text{ for } i \in I\}$ of $\mathcal{WS}_{re}$ and it is a critical point of h restricted to the regular locus of this variety. So $x \in Q_{re} \subset W_e$. □

This leads to Algorithm 2 for computing a finite set containing the limit set $\lim_{r\to 0} \Sigma_{re}$. As an important optimization when the number of equations is large, note that if $\#I \geq n$, then $W'_{I,\sigma,e} = W'_{J,\sigma|_J,e}$ for some $J \subseteq I$ with $\#J = n$. Indeed, if $\mathrm{d}_x h = \sum_{i\in I} \lambda_i \mathrm{d}_x f_i$, then we can find a similar relation using at most n derivatives $\mathrm{d}_x f_i$.

Algorithm 3 Dimension of a real algebraic set (proper case)

Input $f_1, \ldots, f_s \in \mathbb{R}[x_1, \ldots, x_n]$
Precondition The map $x \mapsto (f_1(x), \ldots, f_s(x))$ is proper.
Output Real dimension of $\{x \in \mathbb{R}^n \mid f_1(x) = \cdots = f_s(x) = 0\}$

```
1:  function Dim_proper(f_1, ..., f_s)
2:      if {f_1 = ··· = f_s = 0} = ∅ then
3:          return −1
4:      if n = 1 then
5:          return 0
6:      e ← a generic element of ℝ^s
7:      q ← a generic linear form in x_1, ..., x_{n−1}
8:      Z ← LimitCriticalValues(f_1, ..., f_s, x_n − q, e)
9:      dim ← −1
10:     for U connected component of ℝ \ Z do
11:         t ← some point in U
12:         dimfiber ← Dim_proper(f_1|_{x_n←q+t}, ..., f_s|_{x_n←q+t})
13:         dim ← max(dim, dimfiber + 1)
14:     return dim
```

Proposition 3.5. *On input $f_1, \ldots, f_s$, h and $e \in (\mathbb{R} \setminus \{0\})^s$, and assuming that h is proper on S_e, then Algorithm 2 returns a finite set Z containing $\lim_{r\to 0} \Sigma_{re}$.*

3.2.3 Complexity. When h is generic, the computation of limits of critical values can be performed in the framework of geometric resolution, which leads to complexity bounds. For brevity, we study the case $s = 1$; without much loss of generality since we can replace several equations with a sum of squares.

Proposition 3.6 ([29]). *On input $f \in \mathbb{R}[x_1, \ldots, x_n]$, $h \in \mathcal{L}$ or Q generic and $e \in \mathbb{R}^s$ generic, one can compute a finite set $Z \subset \mathbb{R}$ with less than D^n elements such that $\lim_{r\to 0} \Sigma_{re} \subseteq Z$ in at most $\mathrm{poly}(\log D, n) D^{2n+2}$ arithmetic operations, where $D = \deg f$ and L is the evaluation complexity of f.*

3.3 Computation of the dimension

Theorem 3.7. *On input $f_1, \ldots, f_s$, and assuming that the map $x \in \mathbb{R}^n \mapsto (f_1(x), \ldots, f_s(x))$ is proper, Algorithm 3 generically*[1] *returns the dimension of the real algebraic set $\{f_1 = \cdots = f_s = 0\}$*

Proof. We proceed by induction on n. The case $n = 1$ is trivial. If $V = \{f_1 = \cdots = f_s = 0\}$ is empty, then the algorithm returns -1 on line 3. Assume now that V is not empty. Using Proposition 3.5 and Theorem 2.5, the function $\dim V(t)$ is locally constant on $\mathbb{R} \setminus Z$. So the algorithm computes and return $\max_{t\in\mathbb{R}\setminus Z} \dim V(t)$. By Proposition 2.3, this is $\dim V$. □

The nonproper case is similar.

Theorem 3.8. *On input $f_1, \ldots, f_s$, Algorithm 4 generically returns the dimension of the real algebraic set $\{f_1 = \cdots = f_s = 0\}$.*

We study the complexity in the proper case in the case $s = 1$ in the framework of geometric resolution.

Theorem 3.9. *Given $f \in \mathbb{R}[x_1, \ldots, x_n]$ (of degree D and evaluation complexity L) such that $\mathbb{R}^n \cap \{f = 0\}$ is bounded. One can*

[1] That is, assuming that the points picked at lines 6 and 7 are generic enough.

Algorithm 4 Dimension of a real algebraic set

Input $f_1, \dots, f_s \in \mathbb{R}[x_1, \dots, x_n]$
Output Real dimension of $\{p \in \mathbb{R}^n \mid f_1(p) = \cdots = f_s(p) = 0\}$

1: **function** Dimension($f_1, \dots, f_s$)
2: **if** $\{f_1 = \cdots = f_s = 0\} = \varnothing$ **then**
3: **return** -1
4: $e \leftarrow$ a generic element of $\mathbb{R}^s$
5: $p \leftarrow$ a generic element of $\mathbb{R}^n$
6: $h \leftarrow (x_1 - p_1)^2 + \cdots + (x_n - p_n)^2$
7: $Z \leftarrow$ LimitCriticalValues($f_1, \dots, f_s, h, e$)
8: $\dim \leftarrow -1$
9: **for** U connected component of $\mathbb{R} \setminus Z$ **do**
10: $t \leftarrow$ some point in U
11: $\dim \leftarrow \max(\dim, \mathrm{Dim}_{\mathrm{proper}}(f_1, \dots, f_s, h - t) + 1)$
12: **return** dim

compute the dimension of $\mathbb{R}^n \cap \{f = 0\}$ *in* $\mathrm{poly}(\log D, n) L D^{n(d+3)+1}$ *arithmetic operation and at most* $D^{n(d+1)}$ *isolation of the real roots of polynomials of degree at most* D^n*, where* d *the dimension to be computed.*

Proof. First note that the emptiness test (line 2) can be performed with $\mathrm{poly}(\log D, n) D^{2(n+1)}$ arithmetic operations [29], comparable to the cost of computing the limits of critical values. Let $a_{n,d}$ be the maximum number of arithmetic operations performed by the algorithm given as input s equations of degree D defining an algebraic set of dimension d, excluding the root isolation necessary to idenitfy points in the connected components of $\mathbb{R} \setminus Z$ (line 10). We have $a_{n,-1} = \mathrm{poly}(\log D, n) D^{2(n+1)}$, and given that $\#Z < D^n$, we have $a_{n,d} \le \mathrm{poly}(\log D, n) D^{2(n+1)} + D^n a_{n-1,d-1}$. It follows that

$$\begin{aligned} a_{n,d} &= L \sum_{k=0}^{d+1} D^{n+\cdots+(n-k+1)} \,\mathrm{poly}(\log D, n) D^{2(n-k)+2} \\ &\le \mathrm{poly}(\log D, n) L D^{n(d+3)+1}. \end{aligned}$$ □

3.4 Computation of the dimension (Las Vegas)

In Algorithm 3, the genericity of e may be checked with Algorithm 1. However, we do not know how to check the genericity of h. The problem can be circumvented by considering n linearly independent linear forms $h_1, \dots, h_n$, based on the following statement.

Proposition 3.10. *Let* $V_i(t) = V \cap \{x_i = t\}$*. Let* $Z_1, \dots, Z_n$ *be finite sets such that* $t \mapsto \dim V_i(t)$ *is locally constant on* $\mathbb{R} \setminus Z_i$*. If* $V \ne \varnothing$ *then*

$$\dim V = \max_{1 \le i \le n} \max_{t \in \mathbb{R} \setminus Z_i} \dim V_i(t) + 1.$$

Proof. By Proposition 2.2, $\dim V \ge \max_{t \in \mathbb{R} \setminus Z_i} \dim V_i(t) + 1$ for any i. Assume, for contradiction, that $\dim V \ge \max_{t \in \mathbb{R} \setminus Z_i} \dim V_i(t) + 2$ for all i. Let W the set of points where V has dimension $\dim V$. This is a closed semialgebraic set of dimension d [7, Proposition 2.8.12]. Decomposing V with Hardt's triviality theorem, as in the proof of Proposition 2.2, reveals that $W \subseteq \{x \in \mathbb{R}^n \mid x_i \in Z_i\}$, for any i. This implies that $W \subseteq Z_1 \times \cdots \times Z_n$, so W is finite and $\dim V = \dim W = 0$. This contradicts $\dim V \ge 2$. □

Algorithm 5 Dimension of a real algebraic set (Las Vegas)

Input $f_1, \dots, f_s \in \mathbb{R}[x_1, \dots, x_n]$
Precondition The map $x \mapsto (f_1(x), \dots, f_s(x))$ is proper.
Output Real dimension of $\{x \in \mathbb{R}^n \mid f_1(x) = \cdots = f_s(x) = 0\}$

1: **function** $\mathrm{Dim}_{\mathrm{LV}}(f_1, \dots, f_s)$
2: **if** $\{f_1 = \cdots = f_s = 0\} = \varnothing$ **then**
3: **return** -1
4: **if** $n = 1$ **then**
5: **return** 0
6: **repeat**
7: $e \leftarrow$ a generic element of $\mathbb{R}^s$
8: **until** CheckWhitneyStratification($f_1, \dots, f_s, e$)
9: $\dim \leftarrow -1$
10: **for** $i \in \{1, \dots, n\}$ **do**
11: $Z \leftarrow$ LimitCriticalValues($f_1, \dots, f_s, x_i, e$)
12: **for** U connected component of $\mathbb{R} \setminus Z$ **do**
13: $t \leftarrow$ some point in U
14: $\dim \leftarrow \max(\dim, \mathrm{Dim}_{\mathrm{LV}}(f_1|_{x_i \leftarrow t}, \dots, f_s|_{x_i \leftarrow t}) + 1)$
15: **return** dim

4 EXPERIMENTS

Implementation. We have implemented the Monte-Carlo version of our algorithm (Algorithm 4). We rely on implementations based on Gröbner bases for saturating polynomial ideals by others and solving zero-dimensional polynomial systems exactly by computing rational parametrizations of their solution sets. We use implementations of algorithms such as F_4 [15] and Sparse-FGLM [17] for computing Gröbner bases using graded reverse lexicographical orderings (or elimination algorithms) and change the ordering to a lexicographical one in the zero-dimensional case. We rely on the libraries `FGb` [16] for computing Gröbner bases w.r.t. elimination orderings and `msolve` [6] for solving multivariate polynomial systems. The algorithm itself is implemented using the Maple (version 2020) computer algebra system. In Table 1, we report on timings obtained on several instance described below. All computations have been performed sequentially using an Intel Xeon E7-4820 (2.00 GHz) with 1.5 Tb of RAM.

There are a few significant variations between the text and the implementation, concerning Algorithm 2 for the most part. Firstly, instead of computing the limits of critical values, we compute the limits of critical points. From them, we obtain not only the limits of critical values, but we also decide the emptiness or nonemptiness of the input. Secondly, and this only matters when $s > 1$, instead of introducing and eliminating the variables λ_i we directly deal with a formulation in terms of the minors of the Jacobian matrix. Thirdly, and this also matters only when $s > 1$, we ignore σ. Indeed, we observe that whenever Z and Z' are two subsets of $\mathbb{R}$ satisfying the hypothesis of Proposition 2.3, then so does $Z \cap Z'$. By considering the intersection of all $\lim_{r \to 0} \Sigma_{re}$ as the signs of the coefficients e_i run though all possible configurations, it easy to check that we can remove the loop over σ in Algorithm 2 while preserving the correctness of Algorithm 3.

Benchmarks. We use the following instances to evaluate the efficiency of our algorithm and its implementation, always in the case $s = 1$.

	n	D	dim.	#fibers in depth ...0	...1	...2	...3	...4	max. deg.	ours	CAD	RT	BPR	BS
p_4	4	4	3	2	4	1			16	0.4	0.54	**0.1**	-	2.1
p_5	5	4	2	10	10	1			84	38	25	**3.2**	-	42
p_6	6	4	2	10	10	1			292	786	-	**18**	-	452
p_7	7	4	2	10	10	1			940	**9 450**	-	200 071	-	-
p_8	8	4	2	10	10	1			2568	**538 028**	-	-	-	-
b_4	4	8	1	2	1				400	43	-	**0.9**	-	10
b_5	5	10	1	2	1				2100	8 280	-	**16**	-	650
$d_{3,5}$	5	6	3	8	14	12	1		264	**416**	-	-	-	739
$d_{3,6}$	6	6	4	4	8	12	8	1	288	**836**	-	-	-	28 867
$d_{3,7}$	7	6	4	8	8	12	12	1	288	**400**	-	-	-	-
$d_{4,3}$	3	12	1	12	14				756	**334**	-	-	-	1 816
$d_{4,4}$	4	12	2	16	22	14			2328	**31 060**	-	-	-	-
$s_{3,5}$	5	4	2	4	8	6			176	**84**	-	-	-	8 000
$s_{3,6}$	6	4	3	8	16	10	1		296	**2 144**	-	-	-	-
$s_{3,7}$	7	4	4	6	12	12	10	1	448	**6 411**	-	-	-	-
$s_{4,5}$	5	4	1	4	6				192	**108**	-	-	-	4 473
$s_{4,6}$	6	4	2	4	8	6			344	**1 860**	-	-	-	-
$s_{4,7}$	7	4	3	10	14	14	8		344	**35 814**	-	-	-	-
$s_{5,6}$	6	4	1	2	2				688	**2 329**	-	-	-	-
Vor_1	6	8	4	5	8	6	4	1	544	**61 705**	-	-	-	-
Sottile	4	24	2	1	12	18			4 052	**11 338**	-	-	-	-

Table 1. Timings for computing the real dimension of several instances.
Description of the columns: "n", the number of variables; "D", the degree of th input; "# fibers in depth k", the maximum cardinality of Z at depth k of the recursion; "max. deg.", the maximum degree of all zero-dimensional polynomial systems which are solved during the execution of the algorithm; "ours", timings of our algorithm, in seconds; "CAD", timings of Maple's CAD; "RT", timings of Maple's real triangularization, "BPR", timings of our implementation of [5]; "BS", timings of [4]. Symbol '-' means that the computation was stopped after 2 weeks or 100 times the best runtime achieved by another method.

Family p_n. These are the following polynomials of degree 4:

$$p_n = \left(\sum_{i=1}^{n} x_i^2\right)^2 - 4\sum_{i=1}^{n-1} x_{i+1}^2 x_i^2 - 4x_1^2 x_n^2.$$

These polynomials are sums-of-squares and then non-negative over the reals.

Family b_n. These are the following polynomials of degree $2n$ (where n is the number of variables)

$$b_n = \prod_{i=1}^{n}\left(x_i^2 + n - 1\right) - n^{n-2}\left(\sum_{i=1}^{n} x_i\right)^2.$$

These polynomials were introduced in [20] and are known to be non-negative over the reals as well.

Family $s_{c,n}$. These are polynomials which are sums of squares. We denote by $s_{c,n}$ a sum of squares of c quadrics in $\mathbb{Q}[x_1, \ldots, x_n]$. All these polynomials have degree 4.

Family $d_{k,n}$. These polynomials are discriminants of characteristic polynomials of $k \times k$ symmetric linear matrices with entries in $\mathbb{Q}[x_1, \ldots, x_n]$. Such polynomials are known to be sums-of-squares [26]. Hence whenever it is non-empty, their real solution set has dimension less than $n-1$. Further, these polynomials are denoted by $d_{k,n}$ (we take randomly chosen dense linear entries in the matrix); they have total degree $2k$.

Other polynomials. The example "Vor_1" comes from [14]. The example "Sottile", communicated to us by F. Sottile, arises in enumerative geometry.

Results. In Table 1, we report on results obtained by comparing the implementation of our algorithm with implementations of the Cylindrical Algebraic Decomposition in Maple 2020 (command `CylindricalAlgebraicDecompose`), the algorithm of [10] which decomposes semi-algebraic systems into triangular systems from which the dimension can be read (command `RealTriangularize`), as well as our implementations (using Maple 2020 again) of the algorithm of in [5] and , which are both based on quantifier elimination through the critical point method.

From a practical point of view, the worst algorithm is the one of [5] which cannot solve any of the problems in our benchmark suite, despite the fact that it provides the best theoretical complexity $D^{O(d(n-d))}$. The main reason for that is the too large constant hidden by the "big-O" notation which is here in the exponent.

On benchmarks p_n and b_n, our method suffers from the choice of generic quadratic forms or linear forms compute limits of critical points. On these examples, decomposition methods, and especially the real triangular decomposition [10], perform well because the cells are simple. For p_7 and p_8, the theoretical complexity takes over and our algorithm is faster.

On all other instances of our benchmark suite, our algorithm is easer faster than the state-of-the-art software or it can solve problems which were previously out of reach. This is explained by the fact that, as we expected, the number of fibers to be considered for the recursive calls is lower by several orders of magnitude than the exponential bound D^n. Hence, in practice, one observes a behavior which is far from the complexity $D^{O(nd)}$ estimate.

ACKNOWLEDGMENTS

This work is supported by ANRnder grants ANR-19-CE48-0015 (ECARP), ANR-18-CE33-0011 (Sesame) and ANR-19-CE40-0018 (De rerum natura), and the EU H2020 research and innovation programmender grant 813211 (POEMA).

REFERENCES

[1] B. Bank, M. Giusti, J. Heintz, and G. M. Mbakop. "Polar Varieties, Real Equation Solving, and Data Structures: The Hypersurface Case". In: *J. Complex.* 13.1 (1997), pp. 5–27. DOI: 10/bf9q5q.

[2] B. Bank, M. Giusti, J. Heintz, and L. M. Pardo. "Generalized Polar Varieties: Geometry and Algorithms". In: *J. Complexity* 21.4 (2005), pp. 377–412. DOI: 10/cztmxw.

[3] B. Bank, M. Giusti, J. Heintz, M. Safey El Din, and E. Schost. "On the Geometry of Polar Varieties". In: *Appl. Algebra Eng. Commun. Comput.* 21.1 (2010), pp. 33–83. DOI: 10/ch3xnh.

[4] I. Bannwarth and M. Safey El Din. "Probabilistic Algorithm for Computing the Dimension of Real Algebraic Sets". In: *Proceedings of ISSAC'15.* ACM, 2015, pp. 37–44. DOI: 10/ghzkkt.

[5] S. Basu, R. Pollack, and M.-F. Roy. "Computing the Dimension of a Semi-Algebraic Set". In: *J. Math. Sci.* 134.5 (2006), pp. 2346–2353. DOI: 10/bfxhkk.

[6] J. Berthomieu, C. Eder, and M. Safey El Din. "`msolve`: A Library for Solving Polynomial Systems". available at https://msolve.lip6.fr. 2021.

[7] J. Bochnak, M. Coste, and M.-F. Roy. *Real Algebraic Geometry.* Springer, 1998. DOI: 10/d22h.

[8] J. F. Canny. *The Complexity of Robot Motion Planning.* MIT Press, 1988.

[9] J. Capco, M. S. E. Din, and J. Schicho. "Robots, Computer Algebra and Eight Connected Components". In: *Proceedings of ISSAC'20.* ACM, 2020, pp. 62–69. DOI: 10/ghzkm6.

[10] C. Chen, J. H. Davenport, J. P. May, M. Moreno Maza, B. Xia, and R. Xiao. "Triangular decomposition of semi-algebraic systems". In: *J. Symb. Comp.* 49 (2013), pp. 3–26. DOI: 10/fztd62.

[11] G. E. Collins. "Quantifier Elimination for Real Closed Fields by Cylindrical Algebraic Decompostion". In: *Automata Theory and Formal Languages.* Ed. by H. Brakhage. Lecture Notes in Computer Science. Springer, 1975, pp. 134–183. DOI: 10/cn9jvv.

[12] P. Conti and C. Traverso. "Algorithms for the real radical". 1998.

[13] A. H. Durfee. "Neighborhoods of Algebraic Sets". In: *Trans. Amer. Math. Soc.* 276.2 (1983), pp. 517–517. DOI: 10/bx8m78.

[14] H. Everett, D. Lazard, S. Lazard, and M. Safey El Din. "The Voronoi Diagram of Three Lines". In: *Discrete Comput. Geom.* 42.1 (2009), pp. 94–130. DOI: 10/dpsb5d.

[15] J.-C. Faugère. "A new efficient algorithm for computing Gröbner bases (F_4)". In: *J. Pure Appl. Algebra* 139.1-3 (1999), pp. 61–88. DOI: 10/bpq5dx.

[16] J.-C. Faugère. "FGb: A Library for Computing Gröbner Bases". In: *Mathematical Software – ICMS 2010.* Ed. by K. Fukuda, J. van der Hoeven, M. Joswig, and N. Takayama. Lecture Notes in Computer Science. Springer, 2010, pp. 84–87. DOI: 10/ckcqx9.

[17] J.-C. Faugère and C. Mou. "Sparse FGLM Algorithms". In: *J. Symb. Comput.* 80 (2017), pp. 538–569. DOI: 10/gfz47c.

[18] C. G. Gibson, K. Wirthmüller, A. A. D. Plessis, and E. J. N. Looijenga. *Topological Stability of Smooth Mappings.* Lecture Notes in Mathematics. Springer-Verlag, 1976. DOI: 10/bt4gr4.

[19] M. Giusti, G. Lecerf, and B. Salvy. "A Gröbner Free Alternative for Polynomial System Solving". In: *J. Complexity* 17.1 (2001), pp. 154–211. DOI: 10/fpzjtc.

[20] J. Han, L. Dai, and B. Xia. "Constructing Fewer Open Cells by GCD Computation in CAD Projection". In: *Proceedings of ISSAC'14.* ACM, 2014, pp. 240–247. DOI: 10/ghzknv.

[21] R. M. Hardt. "Semi-Algebraic Local-Triviality in Semi-Algebraic Mappings". In: *Am. J. Math.* 102.2 (1980), p. 291. DOI: 10/b9m293.

[22] K. Harris, J. D. Hauenstein, and A. Szanto. "Smooth Points on Semi-Algebraic Sets". 2020. arXiv: 2002.04707.

[23] A. Hatcher. *Algebraic Topology.* Cambridge University Press, 2002.

[24] P. Koiran. "Randomized and Deterministic Algorithms for the Dimension of Algebraic Varieties". In: *Proceedings of FOCS'97.* IEEE, 1997, pp. 36–45. DOI: 10/c8k78z.

[25] P. Koiran. "The Real Dimension Problem Is $NP_{\mathbb{R}}$-Complete". In: *J. Complex.* 15.2 (1999), pp. 227–238. DOI: 10/fs6pwc.

[26] P. D. Lax. "On the Discriminant of Real Symmetric Matrices". In: *Commun. Pure Appl. Math.* 51.11-12 (1998), pp. 1387–1396. DOI: 10/bc9g35.

[27] J. Mather. "Notes on Topological Stability". In: *Bull. Amer. Math. Soc.* 49.4 (2012), pp. 475–506. DOI: 10/gg2nbz.

[28] M. Mezzarobba and M. Safey El Din. "Computing roadmaps in smooth real algebraic sets". In: *Proceedings of Transgressive Computing*. Ed. by J.-G. Dumas. 2006, pp. 327–338.

[29] M. Safey El Din. *Computing Sampling Points on a Singular Real Hypersurface Using Lagrange's System.* Research Report 5464. INRIA, 2005.

[30] M. Safey El Din and É. Schost. "Polar Varieties and Computation of One Point in Each Connected Component of a Smooth Algebraic Set". In: *Proceedings of ISSAC'03.* ACM, New York, 2003, pp. 224–231. DOI: 10/cp5qsq.

[31] M. Safey El Din, Z.-H. Yang, and L. Zhi. "Computing Real Radicals and S-Radicals of Polynomial Systems". In: *J. Symb. Comput.* 102 (2021), pp. 259–278. DOI: 10/ghzkn2.

[32] N. Vorobjov. "Complexity of Computing the Local Dimension of a Semialgebraic Set". In: *J. Symb. Comput.* 27.6 (1999), pp. 565–579. DOI: 10/fsj467.