

POLYNOMIAL HIERARCHY, BETTI NUMBERS AND A REAL ANALOGUE OF TODA'S THEOREM

SAUGATA BASU AND THIERRY ZELL

ABSTRACT. Toda [35] proved in 1989 that the (discrete) polynomial time hierarchy, \mathbf{PH} , is contained in the class $\mathbf{P}^{\#\mathbf{P}}$, namely the class of languages that can be decided by a Turing machine in polynomial time given access to an oracle with the power to compute a function in the counting complexity class $\#\mathbf{P}$. This result which illustrates the power of counting is considered to be a seminal result in computational complexity theory. An analogous result in the complexity theory over the reals (in the sense of Blum-Shub-Smale real Turing machines [9]) has been missing so far. In this paper we formulate and prove a real analogue of Toda's theorem. Unlike Toda's proof in the discrete case, which relied on sophisticated combinatorial arguments, our proof is topological in nature. As a consequence of our techniques we are also able to relate the computational hardness of two extremely well-studied problems in algorithmic semi-algebraic geometry – namely the problem of deciding sentences in the first order theory of the reals with a constant number of quantifier alternations, and that of computing Betti numbers of semi-algebraic sets. We obtain a polynomial time reduction of the compact version of the first problem to the second. This latter result might be of independent interest to researchers in algorithmic semi-algebraic geometry.

1. INTRODUCTION AND MAIN RESULTS

1.1. History and Background. In this paper we study the relationship between the computational hardness of two important classes of problems in algorithmic semi-algebraic geometry. Algorithmic semi-algebraic geometry is concerned with designing efficient algorithms for deciding geometric as well as topological properties of semi-algebraic sets. There is a large body of research in this area (see [3] for background). If we consider the most important algorithmic problems studied in this area (see for instance the survey article [2]), it is possible to classify them into two broad sub-classes. The first class consists of the problem of quantifier elimination, and its special cases such as deciding a sentence in the first order theory of the reals, or deciding emptiness of semi-algebraic sets (also often called the existential

Date: **March 31, 2009.**

1991 Mathematics Subject Classification. Primary 14P10, 14P25; Secondary 68W30.

Key words and phrases. Polynomial hierarchy, Betti numbers, Semi-algebraic sets, Toda's theorem.

The first author was supported in part by an NSF grant CCF-0634907.

theory of the reals). The existence of algorithms for solving these problems was first proved by Tarski [34] and later research has aimed at designing algorithms with better complexities [28, 23, 22, 5].

The second class of problems in algorithmic semi-algebraic geometry that has been widely investigated consists of computing topological invariants of semi-algebraic sets, such as counting the number of connected components, computing the Euler-Poincaré characteristic, and more generally all the Betti numbers of semi-algebraic sets [13, 24, 21, 1, 7, 4]. Note that the properties such as connectivity or the vanishing of some Betti number of a semi-algebraic set is not expressible in first-order logic, and thus the existence of algorithms for deciding such properties, is not an immediate consequence of Tarski's result but usually requires some additional topological ingredients such as semi-algebraic triangulations or Morse theory etc. Even though the most efficient algorithms for computing the Betti numbers of a semi-algebraic set uses efficient algorithms for quantifier elimination in an essential way [4, 6], the exact relationship between these two classes of problems has not been clarified from the point of view of computational complexity and doing so is one of the motivations of this paper.

The primary motivation for this paper comes from classical (i.e. discrete) computational complexity theory. In classical complexity theory, there is a seminal result due to Toda [35] linking the complexity of counting with that of deciding sentences with a fixed number of quantifier alternations.

More precisely, Toda's theorem gives the following inclusion (see Section 1.2 below or refer to [27] for precise definitions of the complexity classes appearing in the theorem).

Theorem 1.1 (Toda [35]).

$$\mathbf{PH} \subset \mathbf{P}^{\#\mathbf{P}}.$$

In other words, any language in the (discrete) polynomial hierarchy can be decided by a Turing machine in polynomial time, given access to an oracle with the power to compute a function in $\#\mathbf{P}$.

Remark 1.2. The proof of Theorem 1.1 in [35] is quite non-trivial. While it is obvious that the classes \mathbf{P} , \mathbf{NP} , \mathbf{coNP} are contained in $\mathbf{P}^{\#\mathbf{P}}$, the proof for the higher levels of the polynomial hierarchy is quite intricate and proceeds in two steps: first proving that the $\mathbf{PH} \subset \mathbf{BP} \cdot \oplus \cdot \mathbf{P}$ (using previous results of Schöning [29], and Valiant and Vazirani [36]), and then showing that $\mathbf{BP} \cdot \oplus \cdot \mathbf{P} \subset \mathbf{P}^{\#\mathbf{P}}$. Aside from the obvious question about what should be a proper analogue of the complexity class $\#\mathbf{P}$ over the reals, because of the presence of the intermediate complexity class in the proof, there seems to be no direct way of extending such a proof to real complexity classes in the sense of Blum-Shub-Smale model of computation [9, 30]. The proof of the main theorem (Theorem 1.16) of this paper, which can be seen as a real analogue of Theorem 1.1, proceeds along completely different lines and is mainly topological in nature.

In the late eighties Blum, Shub and Smale [9, 30] introduced the notion of Turing machines over more general fields, thereby generalizing the classical problems of computational complexity theory such as \mathbf{P} vs \mathbf{NP} to corresponding problems over arbitrary fields (such as the real, complex, p -adic numbers etc.) If one considers languages accepted by a Blum-Shub-Smale machine over a finite field one recovers the classical notions of discrete complexity theory. Over the last two decades there has been a lot of research activity towards proving real as well as complex analogues of well known theorems in discrete complexity theory. The first steps in this direction were taken by the authors Blum, Shub, and Smale (henceforth B-S-S) themselves, when they proved the \mathbf{NP}_R -completeness of the problem of deciding whether a real polynomial equation in many variables of degree at most four has a real solution (this is the real analogue of Cook-Levin's theorem that the satisfiability problem is \mathbf{NP} -complete in the discrete case), and subsequently through the work of several researchers (Koiran, Bürgisser, Cucker, Meer to name a few) a well-established complexity theory over the reals as well as complex numbers have been built up, which mirrors closely the discrete case.

From the point of view of computational complexity theory of real B-S-S machines the classes \mathbf{PH} and $\#\mathbf{P}$ appearing in the two sides of the inclusion in Theorem 1.1 can be identified with the two broad classes of problems in algorithmic semi-algebraic geometry discussed previously, viz. the polynomial hierarchy with the problem of deciding sentences with a fixed number of quantifier alternations, and the class $\#\mathbf{P}$ with the problem of computing certain topological invariants of semi-algebraic sets, namely their Betti numbers which generalize the notion of cardinality for finite sets. (This naive intuition is made more precise in Section 1.2.2.) It is thus quite natural to seek a real analogue of Toda's theorem. Indeed, there has been a large body of recent research on obtaining appropriate real (as well as complex) analogues of results in discrete complexity theory, especially those related to counting complexity classes (see [26, 10, 12, 11]).

In order to formulate such a result it is first necessary to define precisely real counter-parts of the discrete polynomial time hierarchy \mathbf{PH} and the discrete complexity class $\#\mathbf{P}$, and this is what we do next.

1.2. Real counter-parts of \mathbf{PH} and $\#\mathbf{P}$. For the rest of the paper R will denote a real closed field (there is no essential loss in assuming that $R = \mathbb{R}$). By a real Turing machine we will mean a machine in the sense of Blum-Shub-Smale [9]) over the ground field R .

Notational convention. Since in what follows we will be forced to deal with multiple blocks of variables in our formulas, we follow a notational convention by which we denote blocks of variables by bold letters with superscripts (e.g. \mathbf{X}^i denotes the i -th block), and we use non-bold letters with subscripts to denote single variables (e.g. X_j^i denotes the j -th variable in the i -th block). We use \mathbf{x}^i to denote a specific value of the block of variables \mathbf{X}^i .

1.2.1. *Real analogue of **PH**.* We recall the definition of the polynomial hierarchy for the reals. It mirrors the discrete case very closely (see [33]).

Definition 1.3 (The class \mathbf{P}_R). Let $k(n)$ be any polynomial in n . A sequence of semi-algebraic sets $(T_n \subset \mathbb{R}^{k(n)})_{n>0}$ is said to belong to the class \mathbf{P}_R if there exists a Turing machine M over \mathbb{R} (see [9, 8]), such that for all $\mathbf{x} \in \mathbb{R}^{k(n)}$, the machine M tests membership of \mathbf{x} in T_n in time bounded by a polynomial in n .

Definition 1.4. Let $k(n), k_1(n), \dots, k_\omega(n)$ be polynomials in n . A sequence of semi-algebraic sets $(S_n \subset \mathbb{R}^{k(n)})_{n>0}$ is said to be in the complexity class $\Sigma_{R,\omega}$, if for each $n > 0$ the semi-algebraic set S_n is described by a first order formula

$$(1.1) \quad (Q_1 \mathbf{Y}^1) \cdots (Q_\omega \mathbf{Y}^\omega) \phi_n(X_1, \dots, X_{k(n)}, \mathbf{Y}^1, \dots, \mathbf{Y}^\omega),$$

with ϕ_n a quantifier free formula in the first order theory of the reals, and for each $i, 1 \leq i \leq \omega$, $\mathbf{Y}^i = (Y_1^i, \dots, Y_{k_i(n)}^i)$ is a block of $k_i(n)$ variables, $Q_i \in \{\exists, \forall\}$, with $Q_j \neq Q_{j+1}, 1 \leq j < \omega$, $Q_1 = \exists$, and the sequence of semi-algebraic sets $(T_n \subset \mathbb{R}^{k(n)+k_1(n)+\dots+k_\omega(n)})_{n>0}$ defined by the quantifier-free formulas $(\phi_n)_{n>0}$ belongs to the class \mathbf{P}_R .

Similarly, the complexity class $\Pi_{R,\omega}$ is defined as in Definition 1.4, with the exception that the alternating quantifiers in (1.1) start with $Q_1 = \forall$. Since, adding an additional block of quantifiers on the outside (with new variables) does not change the set defined by a quantified formula we have the following inclusions:

$$\Sigma_{R,\omega} \subset \Pi_{R,\omega+1}, \text{ and } \Pi_{R,\omega} \subset \Sigma_{R,\omega+1}.$$

Note that by the above definition the class $\Sigma_{R,0} = \Pi_{R,0}$ is the familiar class \mathbf{P}_R , the class $\Sigma_{R,1} = \mathbf{NP}_R$ and the class $\Pi_{R,1} = \mathbf{co-NP}_R$.

Definition 1.5 (Real polynomial hierarchy). The real polynomial time hierarchy is defined to be the union

$$\mathbf{PH}_R \stackrel{\text{def}}{=} \bigcup_{\omega \geq 0} (\Sigma_{R,\omega} \cup \Pi_{R,\omega}) = \bigcup_{\omega \geq 0} \Sigma_{R,\omega} = \bigcup_{\omega \geq 0} \Pi_{R,\omega}.$$

For technical reasons (see Remark 2.13) we need to restrict to compact semi-algebraic sets, and for this purpose, we will now define a compact analogue of \mathbf{PH}_R that we will denote \mathbf{PH}_R^c .

Definition 1.6. We call $K \subset \mathbb{R}^n$ a *semi-algebraic compact* if it is a closed and bounded semi-algebraic set. (Note that if $\mathbb{R} \neq \mathbb{R}$, K is not necessarily compact in the order topology.)

Notation 1.7. We denote by $\mathbf{B}^k(0, r)$ the closed ball in \mathbb{R}^k of radius r centered at the origin. We will denote by \mathbf{B}^k the closed unit ball $\mathbf{B}^k(0, 1)$. Similarly, we denote by $\mathbf{S}^k(0, r)$ the sphere in \mathbb{R}^{k+1} of radius r centered at the origin, and by \mathbf{S}^k the unit sphere $\mathbf{S}^k(0, 1)$.

We now define our compact analogue \mathbf{PH}_R^c of the real polynomial hierarchy \mathbf{PH}_R . Unlike in the non-compact case, we will assume all variables vary over certain compact semi-algebraic sets (namely spheres of varying dimensions).

Definition 1.8 (Compact real polynomial hierarchy). Let

$$k(n), k_1(n), \dots, k_\omega(n)$$

be polynomials in n . A sequence of semi-algebraic sets $(S_n \subset \mathbf{S}^{k(n)})_{n>0}$ is in the complexity class $\Sigma_{R,\omega}^c$, if for each $n > 0$ the semi-algebraic set S_n is described by a first order formula

$$(Q_1 \mathbf{Y}^1 \in \mathbf{S}^{k_1(n)}) \dots (Q_\omega \mathbf{Y}^\omega \in \mathbf{S}^{k_\omega(n)}) \phi_n(X_0, \dots, X_{k(n)}, \mathbf{Y}^1, \dots, \mathbf{Y}^\omega),$$

with ϕ_n a quantifier-free first order formula defining a *closed* semi-algebraic subset of $\mathbf{S}^{k_1(n)} \times \dots \times \mathbf{S}^{k_\omega(n)} \times \mathbf{S}^{k(n)}$ and for each $i, 1 \leq i \leq \omega$, $\mathbf{Y}^i = (Y_0^i, \dots, Y_{k_i(n)}^i)$ is a block of $k_i(n) + 1$ variables, $Q_i \in \{\exists, \forall\}$, with $Q_j \neq Q_{j+1}, 1 \leq j < \omega$, $Q_1 = \exists$, and the sequence of semi-algebraic sets $(T_n \subset \mathbf{S}^{k_1(n)} \times \dots \times \mathbf{S}^{k_\omega(n)} \times \mathbf{S}^{k(n)})_{n>0}$ defined by the formulas $(\phi_n)_{n>0}$ belongs to the class \mathbf{P}_R .

Example 1.9. The following is an example of a language in $\Sigma_{R,1}^c$ (i.e. the compact version of \mathbf{NP}_R).

Let $k(n) = \binom{n+4}{4} - 1$ and identify $\mathbf{R}^{k(n)+1}$ with the space of *homogeneous* polynomials in $\mathbf{R}[X_0, \dots, X_n]$ of degree 4. Let $S_n \subset \mathbf{S}^{k(n)} \subset \mathbf{R}^{k(n)+1}$ be defined by

$$S_n = \{P \in \mathbf{S}^{k(n)} \mid \exists \mathbf{x} = (x_0 : \dots : x_n) \in \mathbb{P}_R^n \text{ with } P(\mathbf{x}) = 0\};$$

in other words S_n is the set of (normalized) real forms of degree 4 which have a zero in the real projective space \mathbb{P}_R^n . Then

$$(S_n \subset \mathbf{S}^{k(n)})_{n>0} \in \Sigma_{R,1}^c,$$

since it is easy to see that S_n also admits the description:

$$S_n = \{P \in \mathbf{S}^{k(n)} \mid \exists \mathbf{x} \in \mathbf{S}^n \text{ with } P(\mathbf{x}) = 0\}.$$

Note that it is *not known* if $(S_n \subset \mathbf{S}^{k(n)})_{n>0}$ is \mathbf{NP}_R -complete (see Remark 1.10), while the non-compact version of this language i.e. the language consisting of (possibly non-homogeneous) polynomials of degree at most four having a zero in \mathbb{A}_R^n (instead of \mathbb{P}_R^n), has been shown to be \mathbf{NP}_R -complete [8].

We define analogously the class $\Pi_{R,\omega}^c$, and finally define the *compact real polynomial time hierarchy* to be the union

$$\mathbf{PH}_R^c \stackrel{\text{def}}{=} \bigcup_{\omega \geq 0} (\Sigma_{R,\omega}^c \cup \Pi_{R,\omega}^c) = \bigcup_{\omega \geq 0} \Sigma_{R,\omega}^c = \bigcup_{\omega \geq 0} \Pi_{R,\omega}^c.$$

Notice that the semi-algebraic sets belonging to any language in \mathbf{PH}_R^c are all semi-algebraic compact (in fact closed semi-algebraic subsets of spheres). Also, note the inclusion

$$\mathbf{PH}_R^c \subset \mathbf{PH}_R.$$

Remark 1.10. Even though the restriction to compact semi-algebraic sets might appear to be only a technicality at first glance, this is actually an important restriction. For instance, it is a long-standing open question in real complexity theory whether there exists an \mathbf{NP}_R -complete problem which belongs to the class $\Sigma_{R,1}^c$ (the compact version of the class \mathbf{NP}_R , see Example 1.9). (This distinction between compact and non-compact versions of complexity classes does not arise in discrete complexity theory for obvious reasons.) It is an interesting question whether the main theorem of this paper can be extended to the full class \mathbf{PH}_R . For technical reasons which will become clear later in the paper (Remark 2.13) we are unable to achieve this presently.

1.2.2. *Real Analogue of $\#\mathbf{P}$.* Before defining the real analogue of the class $\#\mathbf{P}$, let us recall its definition in the discrete case which is well known.

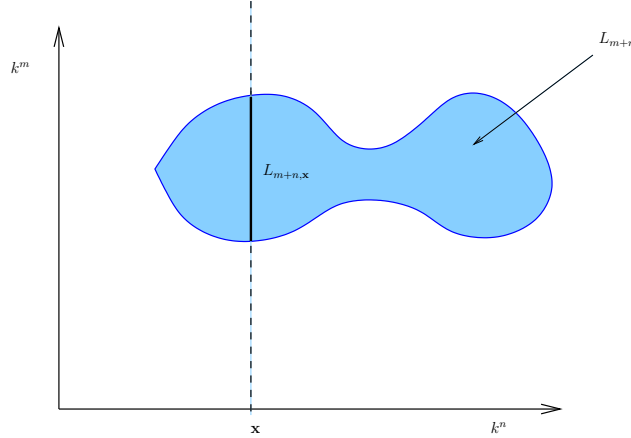


FIGURE 1. The fibers $L_{m+n,\mathbf{x}}$ of the language L

Definition 1.11. Let $k = \mathbb{Z}/2\mathbb{Z}$. We say that a sequence of functions

$$(f_n : k^n \rightarrow \mathbb{N})_{n>0}$$

is in the class $\#\mathbf{P}$ if there exists a language

$$L = \bigcup_{n>0} L_n, \quad L_n \subset k^n, \quad \text{with } L \in \mathbf{P}$$

as well as a polynomial $m(n)$, such that

$$f_n(\mathbf{x}) = \text{card}(L_{m+n,\mathbf{x}})$$

for each $\mathbf{x} \in k^n$, where $L_{m+n,\mathbf{x}} = L_{m+n} \cap \pi^{-1}(\mathbf{x})$ and $\pi : k^{m+n} \rightarrow k^n$ is the projection along the first m co-ordinates.

In other words, f_n counts the number of points in the fibers, $L_{m+n,\mathbf{x}}$, of a language L belonging to the (discrete) complexity class \mathbf{P} (see Figure 1). (The geometric language used above might look unnecessary but it is very helpful towards obtaining the right analogue in the real case.)

In order to define real analogues of counting complexity classes of discrete complexity theory, it is necessary to identify the proper notion of “counting” in the context of semi-algebraic geometry. Counting complexity classes over the reals have been defined previously by Meer [26], and studied extensively by other authors [11]. These authors used a straightforward generalization to semi-algebraic sets of counting in the case of finite sets – namely the counting function took the value of the cardinality of a semi-algebraic set if it happened to be finite, and ∞ otherwise. This is in our view not a fully satisfactory generalization since the count gives no information when the semi-algebraic set is infinite, and most interesting semi-algebraic sets have infinite cardinality. Moreover, no real analogue of Toda’s theorem has been proved using this definition of counting.

If one thinks of “counting” a semi-algebraic set $S \subset \mathbb{R}^k$ as computing a certain discrete invariant, then a natural well-studied discrete topological invariant of S is its Euler-Poincaré characteristic. For a closed and bounded semi-algebraic set S the Euler-Poincaré characteristic, $\chi(S)$, is the alternating sum of the Betti numbers of S , and it is possible to extend this definition to the class of all semi-algebraic sets by additivity. This generalized Euler-Poincaré characteristic gives an isomorphism from the *Grothendieck ring of semi-algebraic sets* to \mathbb{Z} , and thus corresponds to a mathematically natural notion of counting semi-algebraic sets. However, the Euler-Poincaré characteristic fails to distinguish between empty and non-empty semi-algebraic sets, since a non-empty semi-algebraic set (e.g. an odd dimensional sphere) can have vanishing Euler-Poincaré characteristic. This seems to immediately rule out using the Euler-Poincaré characteristic as a substitute for the counting function. We make up for this deficiency by replacing the Euler-Poincaré characteristic by the Poincaré polynomial $P_S(T)$ of the set S . We now recall the relevant definitions.

Notation 1.12. For any semi-algebraic set $S \subset \mathbb{R}^k$ we denote by $b_i(S)$ the i -th Betti number (that is the rank of the singular homology group $H_i(S) = H_i(S, \mathbb{Z})$) of S .

We also let $P_S \in \mathbb{Z}[T]$ denote the *Poincaré polynomial* of S , namely

$$(1.2) \quad P_S(T) \stackrel{\text{def}}{=} \sum_{i \geq 0} b_i(S) T^i.$$

Notice that for $S \subset \mathbb{R}^k$, $\deg(P_S) \leq k - 1$. Also, it is easy to see that the Poincaré polynomial, $P_S(T)$, carries more complete information

about S than its Euler-Poincaré characteristic. Indeed, the number of semi-algebraically connected components, $b_0(S)$, of S is obtained by setting T to 0, and in case S is closed and bounded we also recover $\chi(S)$ by setting T to -1 in $P_S(T)$. Since $b_0(S) > 0$ if and only if S is non-empty, P_S , unlike $\chi(S)$, can distinguish between empty and non-empty semi-algebraic sets. In particular, in case S is a finite set of points, P_S also contains the information regarding the cardinality of S which in this case equals $b_0(S) = P_S(0)$.

Remark 1.13. The connection between counting points of varieties and their Betti numbers is more direct over fields of positive characteristic via the zeta function. The zeta function of a variety defined over \mathbb{F}_p is the exponential generating function of the sequence whose n -th term is the number of points in the variety over \mathbb{F}_{p^n} . The zeta function of such a variety turns out to be a rational function in one variable (a deep theorem of algebraic geometry first conjectured by Andre Weil [37] and proved by Dwork [17] and Deligne [15, 16]), and its numerator and denominator are products of polynomials whose degrees are the Betti numbers of the variety with respect to a certain (ℓ -adic) co-homology theory. The point of this remark is that the problems of “counting” varieties and computing their Betti numbers, are connected at a deeper level, and thus our choice of definition for a real analogue of $\#\mathbf{P}$ is not merely ad hoc.

The above considerations motivate us to depart from the definition of $\#\mathbf{P}_R$ considered previously in [26, 11]. We denote our class $\#\mathbf{P}_R^\dagger$ to avoid any possible confusion with these authors’ work.

Definition 1.14 (The class $\#\mathbf{P}_R^\dagger$). We say a sequence of functions

$$(f_n : \mathbb{R}^n \rightarrow \mathbb{Z}[T])_{n>0}$$

is in the class $\#\mathbf{P}_R^\dagger$, if there exists a language

$$(S_n \subset \mathbb{R}^n)_{n>0} \in \mathbf{P}_R,$$

as well as a polynomial $m(n)$, such that

$$f_n(\mathbf{x}) = P_{S_{m+n}, \mathbf{x}}$$

for each $\mathbf{x} \in \mathbb{R}^n$, where $S_{m+n, \mathbf{x}} = S_{m+n} \cap \pi^{-1}(\mathbf{x})$ and $\pi : \mathbb{R}^{m+n} \rightarrow \mathbb{R}^n$ is the projection along the first m co-ordinates.

Remark 1.15. Notice the formal similarity between Definitions 1.14 and 1.11, namely that in both cases the functions f_n “counts” the fibers above \mathbf{x} , but the notion of counting is different in each case.

1.3. Statements of the main theorems. We can now state the main result of this paper.

Theorem 1.16 (Real analogue of Toda’s theorem).

$$\mathbf{PH}_R^c \subset \mathbf{P}_R^{\#\mathbf{P}_R^\dagger}.$$

Remark 1.17. We leave it as an open problem to prove Theorem 1.16 with \mathbf{PH}_R instead of \mathbf{PH}_R^c on the left hand side. One possible approach would be to use the recent results of Gabrielov and Vorobjov [19] on replacing arbitrary semi-algebraic sets by compact semi-algebraic sets in the same homotopy equivalence class using infinitesimal deformations. However, for such a construction to be useful in our context, one would need to effectively (i.e. in polynomial time) replace the infinitesimals used in the construction by small enough positive elements of R , and at present we are unable to achieve this.

As a consequence of our method, we obtain a reduction (Theorem 1.20) that might be of independent interest. We first define the following two problems:

Definition 1.18. (Compact general decision problem with at most ω quantifier alternations (\mathbf{GDP}_ω^c))

Input. A sentence Φ in the first order theory of R

$$(Q_1 \mathbf{X}^1 \in \mathbf{S}^{k_1}) \cdots (Q_\omega \mathbf{X}^\omega \in \mathbf{S}^{k_\omega}) \phi(\mathbf{X}^1, \dots, \mathbf{X}^\omega),$$

where for each $i, 1 \leq i \leq \omega$, $\mathbf{X}^i = (X_0^i, \dots, X_{k_i}^i)$ is a block of $k_i + 1$ variables, $Q_i \in \{\exists, \forall\}$, with $Q_j \neq Q_{j+1}, 1 \leq j < \omega$, and ϕ is a quantifier-free formula defining a *closed* semi-algebraic subset S of $\mathbf{S}^{k_1} \times \dots \times \mathbf{S}^{k_\omega}$.

Output. True or False depending on whether Φ is true or false in the first order theory of R .

Definition 1.19. (Computing the Poincaré polynomial of semi-algebraic sets ($\mathbf{Poincaré}$))

Input. A quantifier-free formula defining a semi-algebraic set $S \subset R^k$.

Output. The Poincaré polynomial $P_S(T)$.

Theorem 1.20. *For every $\omega > 0$, there is a deterministic polynomial time reduction in the Blum-Shub-Smale model of \mathbf{GDP}_ω^c to $\mathbf{Poincaré}$.*

1.4. Summary of the main ideas. Our main tool is a topological construction described in Section 3, which, given a semi-algebraic set $S \subset R^{m+n}$, $p \geq 0$, and $\pi_Y : R^{m+n} \subset R^n$ the projection along (say) the Y co-ordinates, constructs *efficiently* a semi-algebraic set, $D_Y^p(S)$, such that

$$(1.3) \quad b_i(\pi_Y(S)) = b_i(D_Y^p(S)), \quad 0 \leq i < p$$

(in fact, for technical reasons we need two different constructions depending on whether S is an open or a closed semi-algebraic set, but we prefer to ignore this point in this rough outline). An infinitary version of such a construction (and indeed some of the basic ideas behind this construction) is described in [20]. However, the main goal in [20] was to obtain upper bounds on the Betti numbers of semi-algebraic (as well as semi-Pfaffian) sets defined by quantified formulas, and this is achieved by bounding the Betti numbers of certain sets appearing in the E_1 term of a certain (the so

called “descent”) spectral sequence which is guaranteed to converge to the homology of the given set.

In this paper we need to be able to recover exactly (not just bound) the Betti numbers of $\pi_{\mathbf{Y}}(S)$ from those of $D_{\mathbf{Y}}^p(S)$. Moreover, it is very important in our context that membership in the semi-algebraic set $D_{\mathbf{Y}}^p(S)$ should be checkable in polynomial time, given that the same is true for S . Notice that even if there exists an efficient (i.e. polynomial time) algorithm for checking membership in S , the same need not be true for the image $\pi_{\mathbf{Y}}(S)$.

We will now illustrate how the $D_{\mathbf{Y}}^p(S)$ construction connects decision problems in the compact real polynomial hierarchy to the computation of Betti numbers of semi-algebraic sets, by looking at the special case of one and two quantifier alternations.

1.4.1. Case of one quantifier. First consider the class $\Sigma_{\mathbf{R},1}^c$. Consider a closed semi-algebraic set $S \subset \mathbf{S}^k \times \mathbf{S}^\ell$ defined by a quantifier-free formula $\phi(\mathbf{X}, \mathbf{Y})$ and let

$$\pi_{\mathbf{Y}} : \mathbf{S}^k \times \mathbf{S}^\ell \rightarrow \mathbf{S}^k$$

be the projection map along the \mathbf{Y} variables.

Then the formula

$$\Phi(\mathbf{X}) = \exists \mathbf{Y} \phi(\mathbf{X}, \mathbf{Y})$$

is satisfied by $\mathbf{x} \in \mathbf{S}^k$ if and only if $b_0(S_{\mathbf{x}}) \neq 0$, where $S_{\mathbf{x}} = S \cap \pi_{\mathbf{Y}}^{-1}(\mathbf{x})$. Thus, the problem of deciding the truth of $\Phi(\mathbf{x})$ is reduced to computing a Betti number (the 0-th) of the fiber of S over \mathbf{x} .

Now consider the class $\Pi_{\mathbf{R},1}^c$. Using the same notation as above we have that the formula

$$\Psi(\mathbf{X}) = \forall \mathbf{Y} \phi(\mathbf{X}, \mathbf{Y})$$

is satisfied by $\mathbf{x} \in \mathbf{S}^k$ if and only if the formula

$$\neg \Psi(\mathbf{X}) = \exists \mathbf{Y} \neg \phi(\mathbf{X}, \mathbf{Y})$$

does not hold, which means, according to the previous case, that we have $b_0(\mathbf{S}^\ell \setminus S_{\mathbf{x}}) = 0$, which is equivalent to $b_\ell(S_{\mathbf{x}}) = 1$. Notice that, as before, the problem of deciding the truth of $\Psi(\mathbf{x})$ is reduced to computing a Betti number (the ℓ -th) of the fiber of S over \mathbf{x} .

1.4.2. Case of two quantifiers. Proceeding to a slightly more non-trivial case, consider the class $\Pi_{\mathbf{R},2}^c$ and let $S \subset \mathbf{S}^k \times \mathbf{S}^\ell \times \mathbf{S}^m$ be a closed semi-algebraic set defined by a quantifier-free formula $\phi(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ and let

$$\pi_{\mathbf{Z}} : \mathbf{S}^k \times \mathbf{S}^\ell \times \mathbf{S}^m \rightarrow \mathbf{S}^k \times \mathbf{S}^\ell$$

be the projection map along the \mathbf{Z} variables, and

$$\pi_{\mathbf{Y}} : \mathbf{S}^k \times \mathbf{S}^\ell \rightarrow \mathbf{S}^k$$

be the projection map along the \mathbf{Y} variables as before. Consider the formula

$$\Phi(\mathbf{X}) = \forall \mathbf{Y} \exists \mathbf{Z} \phi(\mathbf{X}, \mathbf{Y}, \mathbf{Z}).$$

This formula can be recast as:

$$\Phi(\mathbf{X}) = \forall \mathbf{Y} (\mathbf{X}, \mathbf{Y}) \in \pi_{\mathbf{Z}}(S).$$

Thus, for any $\mathbf{x} \in \mathbf{S}^k$, $\Phi(\mathbf{x})$ holds if and only if we have the following situation:

$$\begin{array}{ccccc} S \hookrightarrow & \mathbf{S}^k \times \mathbf{S}^\ell \times \mathbf{S}^m & & & \\ \downarrow \pi_{\mathbf{Z}} & & \downarrow \pi_{\mathbf{Z}} & & \\ \{\mathbf{x}\} \times \mathbf{S}^\ell \hookrightarrow & \pi_{\mathbf{Z}}(S) \hookrightarrow & \mathbf{S}^k \times \mathbf{S}^\ell & & \\ & \downarrow \pi_{\mathbf{Y}} & \downarrow \pi_{\mathbf{Y}} & & \\ & \mathbf{x} \in \pi_{\mathbf{Y}, \mathbf{Z}}(S) \hookrightarrow & \mathbf{S}^k & & \end{array}$$

i.e. if and only if the $\pi_{\mathbf{Y}}$ fiber $(\pi_{\mathbf{Z}}(S))_{\mathbf{x}}$ is equal to \mathbf{S}^ℓ . This can be formulated in terms of Betti numbers by the condition:

$$b_\ell((\pi_{\mathbf{Z}}(S))_{\mathbf{x}}) = 1.$$

The construction mentioned in (1.3) gives, for $p = \ell + 1$, the existence of a semi-algebraic set $D_{\mathbf{Z}}^{\ell+1}(S)$ such that $b_\ell(D_{\mathbf{Z}}^{\ell+1}(S)) = b_\ell(\pi_{\mathbf{Z}}(S))$. Fortunately, the construction of the set $D_{\mathbf{Z}}^{\ell+1}(S)$ is compatible with taking fibers, so that we have, for all $\mathbf{x} \in \mathbf{S}^k$,

$$b_\ell((\pi_{\mathbf{Z}}(S))_{\mathbf{x}}) = b_\ell(D_{\mathbf{Z}}^{\ell+1}(S)_{\mathbf{x}}).$$

Thus for any $\mathbf{x} \in \mathbf{S}^k$, the truth or falsity of $\Phi(\mathbf{x})$ is determined by a certain Betti number of the fiber $D_{\mathbf{Z}}^{\ell+1}(S)_{\mathbf{x}}$ over \mathbf{x} of a certain semi-algebraic set $D_{\mathbf{Z}}^{\ell+1}(S)$ which can be constructed efficiently in terms of the set S . The idea behind the proof of the main theorem is a recursive application of the above argument in case when the number of quantifier alternations is larger (but still bounded by some constant) while keeping track of the growth in the sizes of the intermediate formulas and also the number of quantified variables.

The rest of the paper is organized as follows. In Section 2 we fix notation, and prove the topological results needed for the proof of the two main theorems. In Section 3 we describe the semi-algebraic construction of the sets $D_{\mathbf{Y}}^p(S)$ alluded to above and prove its important properties. We prove the main theorems in Section 4.

2. TOPOLOGICAL INGREDIENTS

We first fix a notation.

Notation 2.1. For each $p \geq 0$ we denote

$$\Delta^p = \{(t_0, \dots, t_p) \mid t_i \geq 0, 0 \leq i \leq p, \sum_{i=0}^p t_i = 1\}$$

the standard p -simplex.

We now describe some constructions in algebraic topology which will be useful later in the paper.

2.1. Properties of the join. We first recall the definition of the join of two topological spaces X and Y .

Definition 2.2. The join $J(X, Y)$ of two topological spaces X and Y is defined by

$$(2.1) \quad J(X, Y) \stackrel{\text{def}}{=} X \times Y \times \Delta^1 / \sim,$$

where

$$(x, y, t_0, t_1) \sim (x', y', t_0, t_1)$$

if $t_0 = 1, x = x'$ or $t_1 = 1, y = y'$.

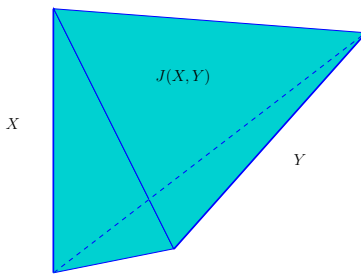


FIGURE 2. Join of two segments

Intuitively, $J(X, Y)$ is obtained by joining each point of X with each point of Y by a unit interval (see Figure 2).

Example 2.3. It is easy to check from the above definition that the join, $J(\mathbf{S}^m, \mathbf{S}^n)$, of two spheres is again (homeomorphic to) a sphere, namely \mathbf{S}^{m+n+1} .

By iterating the above definition with the same space X we obtain

Definition 2.4. For $p \geq 0$ the $(p+1)$ -fold join $J^p(X)$ of X is

$$(2.2) \quad J^p(X) \stackrel{\text{def}}{=} \underbrace{X \times \cdots \times X}_{(p+1) \text{ times}} \times \Delta^p / \sim,$$

where

$$(x_0, \dots, x_p, t_0, \dots, t_p) \sim (x'_0, \dots, x'_p, t_0, \dots, t_p)$$

if for each i with $t_i \neq 0$, $x_i = x'_i$.

Example 2.5. Using Example 2.3 it is easy to see that the $(p+1)$ -fold join, $J^p(\mathbf{S}^0)$, of the zero dimensional sphere is homeomorphic to \mathbf{S}^p .

We will need the fact that the iterated join of a topological space is highly connected. In order to make this statement precise we first define

Definition 2.6 (*p-equivalence*). A map $f : A \rightarrow B$ between two topological spaces is called a *p-equivalence* if the induced homomorphism

$$f_* : H_i(A) \rightarrow H_i(B)$$

is an isomorphism for all $0 \leq i < p$, and an epimorphism for $i = p$, and we say that A is *p-equivalent* to B . (Note that *p-equivalence* is not an equivalence relation : e.g. for any $p \geq 0$, the map taking S^p to a point is a *p-equivalence*, but no map from a point into S^p is one).

Observe from Example 2.5 that $J^p(\mathbf{S}^0) \cong \mathbf{S}^p$ is *p-equivalent* to a point. In fact, this holds much more generally and we have that

Theorem 2.7. *Let X be a compact semi-algebraic set. Then, the $(p+1)$ -fold join $J^p(X)$ is *p-equivalent* to a point.*

Proof. This is classical (see for instance [25, Proposition 4.4.3]). □

2.2. Join over a map. In our application we need the join construction over certain class of maps (to be specified later). We first recall the notion of a fibered product of a topological space.

Notation and definition 2.8. *Let $f : A \rightarrow B$ be a map between topological spaces A and B . For each $p \geq 0$, We denote by $W_f^p(A)$ the $(p+1)$ -fold fiber product of A over f . In other words*

$$W_f^p(A) = \{(x_0, \dots, x_p) \in A^{p+1} \mid f(x_0) = \dots = f(x_p)\}.$$

Definition 2.9 (Topological join over a map). Let $f : A \rightarrow B$ be a map between topological spaces A and B . For $p \geq 0$ the $(p+1)$ -fold join $J_f^p(A)$ of A over f is

$$(2.3) \quad J_f^p(A) \stackrel{\text{def}}{=} W_f^p(A) \times \Delta^p / \sim,$$

where

$$(x_0, \dots, x_p, t_0, \dots, t_p) \sim (x'_0, \dots, x'_p, t_0, \dots, t_p)$$

if for each i with $t_i \neq 0$, $x_i = x'_i$.

We now impose certain conditions on the map f .

2.3. Compact Coverings. Recall that we call $K \subset \mathbb{R}^n$ a semi-algebraic compact if it is a closed and bounded semi-algebraic set.

Notation 2.10. For any semi-algebraic $A \subset \mathbb{R}^n$, we denote by $K(A)$ the collection of all semi-algebraic compact subsets of A .

Definition 2.11. Let $f : A \rightarrow B$ be a semi-algebraic map. We say that f *covers semi-algebraic compacts* if for any $L \in K(f(A))$, there exists $K \in K(A)$ such that $f(K) = L$.

The following theorem relates the topology of $J^p(A)$ to that of the image of f in the case when f covers semi-algebraic compacts and is crucial for what follows.

Theorem 2.12. *Let $f : A \rightarrow B$ be a semi-algebraic map that covers semi-algebraic compacts. Then for every $p \geq 0$, the map f induces a p -equivalence $J(f) : J_f^p(A) \rightarrow f(A)$.*

Proof. We begin with the case $A \in \mathbf{K}(\mathbb{R}^n)$. Let $J(f) : J_f^p(A) \rightarrow f(A)$ be the map given by

$$J(f)(x_0, \dots, x_p, t_0, \dots, t_p) = f(x_0).$$

The map $J(f)$ is well defined since $(x_0, \dots, x_p) \in W_f^p(A)$, and is closed since $J_f^p(A)$ is a semi-algebraic compact. Moreover, the fibers of $J(f)$ are p -equivalent to a point by Theorem 2.7.

Thus, by the Vietoris-Begle theorem [31], the map $J(f)$ induces isomorphisms

$$J(f)_* : H_i(J_f^p(A)) \rightarrow H_i(f(A));$$

for $0 \leq i < p$. Note that in the case $\mathbf{R} \neq \mathbb{R}$, the validity of the Vietoris-Begle theorem can be seen as a corollary of the existence of a semi-algebraic co-homology that satisfies the Eilenberg-Steenrod axioms for a Čech theory (see [18]).

In the general case, consider $K_1 \subset K_2$ two semi-algebraic compacts in $\mathbf{K}(A)$. The inclusion gives rise to the following diagram,

$$\begin{array}{ccc} J_f^p(K_1) & \xhookrightarrow{i} & J_f^p(K_2) \\ \downarrow J(f|_{K_1}) & & \downarrow J(f|_{K_2}) \\ f(K_1) & \xhookrightarrow{j} & f(K_2) \end{array}$$

where the vertical maps are p -equivalence by the previous case. We have a similar diagram at the homology level; if we take the direct limit as K ranges in $\mathbf{K}(A)$, we obtain the following:

$$\begin{array}{ccc} \varinjlim H_i(J_f^p(K)) & \xrightarrow{\cong} & H_i(J_f^p(A)) \\ \downarrow \varinjlim J(f|_K) & & \downarrow J(f) \\ \varinjlim H_i(f(K)) & \xrightarrow{\cong} & H_i(f(A)) \end{array}$$

The isomorphism on the top level comes from the fact that homology and direct limit commute [32], along with the fact that for a semi-algebraic set, one can compute the homology using chains supported exclusively on semi-algebraic compacts [14]. For the bottom isomorphism, we need the additional fact that since we assume that f covers semi-algebraic compacts, we have

$$\varinjlim \{H_i(f(K)) \mid K \in \mathbf{K}(A)\} = \varinjlim \{H_i(L) \mid L \in \mathbf{K}(B)\}.$$

Since each $J(f|_K)$ is a p -equivalence, the vertical homomorphisms are isomorphisms for $0 \leq i < p$, and an epimorphisms for $i = p$. \square

Remark 2.13. Theorem 2.12 requires that the map f covers semi-algebraic compacts. This condition is satisfied for a projection in the case the set A is either open or compact. Note also that Theorem 2.12 is not true without the assumption that f covers semi-algebraic compacts, which is why, in this paper, we restrict our attention to the *compact* polynomial hierarchy.

2.4. Alexander-Lefschetz duality. We will also need the classical Alexander-Lefschetz duality theorem in order to relate the Betti numbers of a compact semi-algebraic subset of a sphere to those of its complement.

Theorem 2.14 (Alexander-Lefschetz duality). *Let $K \subset \mathbf{S}^n$ be a compact semi-algebraic subset with $n \geq 2$. Then*

$$\begin{aligned} b_0(K) &= 1 + b_{n-1}(\mathbf{S}^n - K) - b_n(\mathbf{S}^n - K), \\ b_i(K) &= b_{n-i-1}(\mathbf{S}^n - K), \quad 1 \leq i \leq n-2, \\ b_{n-1}(K) &= b_0(\mathbf{S}^n - K) - 1 + \max(1 - b_0(\mathbf{S}^n - K), 0), \\ b_n(K) &= 1 - \min(1, b_0(\mathbf{S}^n - K)). \end{aligned}$$

Proof. Lefschetz duality theorem [32] gives for each $i, 0 \leq i \leq n$,

$$b_i(\mathbf{S}^n - K) = b_{n-i}(\mathbf{S}^n, K).$$

The theorem now follows from the long exact sequence of homology,

$$\cdots \rightarrow H_i(K) \rightarrow H_i(\mathbf{S}^n) \rightarrow H_i(\mathbf{S}^n, K) \rightarrow H_{i-1}(K) \rightarrow \cdots$$

after noting that $H_i(\mathbf{S}^n) = 0, i \neq 0, n$ and $H_0(\mathbf{S}^n) = H_n(\mathbf{S}^n) = \mathbb{Z}$. \square

3. SEMI-ALGEBRAIC CONSTRUCTIONS

In this section we describe the semi-algebraic construction that lies at the heart of the proof of our main theorem, and prove its important properties.

Let $S \subset \mathbf{S}^k \times \mathbf{S}^\ell$ be a subset defined by a first-order formula $\Phi(\mathbf{X}, \mathbf{Y})$, and let $\pi_{\mathbf{Y}}$ denote the projection along the \mathbf{Y} co-ordinates. Note that \mathbf{X} and \mathbf{Y} are the *free* variables in Φ , the formula being considered may have any number of quantified blocks of variables $(\mathbf{Z}^1, \dots, \mathbf{Z}^\omega)$. This will be addressed explicitly in Lemma 3.6.

We now define semi-algebraic sets having the same homotopy type as the join space $J_{\pi_{\mathbf{Y}}}(S)$ in the case when S is a closed (respectively open) semi-algebraic subset of $\mathbf{S}^k \times \mathbf{S}^\ell$.

Notation and definition 3.1. *Let $S \subset \mathbf{S}^k \times \mathbf{S}^\ell$, $\Phi(\mathbf{X}, \mathbf{Y})$, and $\pi_{\mathbf{Y}}$ be as above. If S is closed, we denote by $D_{\mathbf{Y},c}^p(S)$ the semi-algebraic set defined by*

$$\begin{aligned} D_{\mathbf{Y},c}^p(S) &\stackrel{\text{def}}{=} \{(u, \mathbf{x}, \mathbf{y}^0, \dots, \mathbf{y}^p, \mathbf{t}) \mid \mathbf{x} \in \mathbf{S}^k, \mathbf{t} \in \Delta^p, \\ (3.1) \quad &\text{for each } i, 0 \leq i \leq p, \mathbf{y}^i \in \mathbf{B}^{\ell+1}, (t_i = 0) \vee \Phi(\mathbf{x}, \mathbf{y}^i), \\ &u^2 + |\mathbf{x}|^2 + \sum_{i=0}^p |\mathbf{y}^i|^2 + |\mathbf{t}|^2 = p + 4, \text{ and } u \geq 0\}. \end{aligned}$$

Notice that $D_{\mathbf{Y},c}^p(S)$ is a closed semi-algebraic subset of the upper hemisphere of the sphere $\mathbf{S}^N(0, p+4)$, where $N = (k+1) + (p+1)(\ell+2)$.

We will denote by $D_{\mathbf{Y},c}^p(\Phi)$ the first-order formula defining the semi-algebraic set $D_{\mathbf{Y},c}^p(S)$, namely

$$(3.2) \quad D_{\mathbf{Y},c}^p(\Phi) \stackrel{\text{def}}{=} \Theta_1(T) \wedge \Theta_2(\mathbf{X}, \mathbf{Y}^0, \dots, \mathbf{Y}^p, \mathbf{T}) \wedge \Theta_3(U_0, \mathbf{X}, \mathbf{Y}^0, \dots, \mathbf{Y}^p, \mathbf{T})$$

where

$$\begin{aligned} \Theta_1 &\stackrel{\text{def}}{=} \left(\bigwedge_{i=0}^p T_i \geq 0 \right) \wedge \left(\sum_{i=0}^p T_i = 1 \right), \\ \Theta_2 &\stackrel{\text{def}}{=} \left((|\mathbf{X}|^2 = 1) \bigwedge_{i=0}^p ((|\mathbf{Y}^i|^2 \leq 1) \wedge ((T_i = 0) \vee \Phi(\mathbf{X}, \mathbf{Y}^i))) \right), \\ \Theta_3 &\stackrel{\text{def}}{=} \left(U_0^2 + |\mathbf{X}|^2 + \sum_{i=0}^p |\mathbf{Y}^i|^2 + |\mathbf{T}|^2 = p+4 \right) \wedge (U_0 \geq 0). \end{aligned}$$

We have a similar construction in case S is an open subset of $\mathbf{S}^k \times \mathbf{S}^\ell$. In this case we thicken the various faces of the standard simplex Δ^p (see Figure 3) so that they become convex open subsets of \mathbf{R}^{p+1} , but maintaining the property that a subset of these thickened faces have a non-empty intersection if and only if the closures of the corresponding faces in Δ^p had a non-empty intersection. In this way we ensure that our construction produces an open subset of a sphere, while having again the homotopy type of the join space.

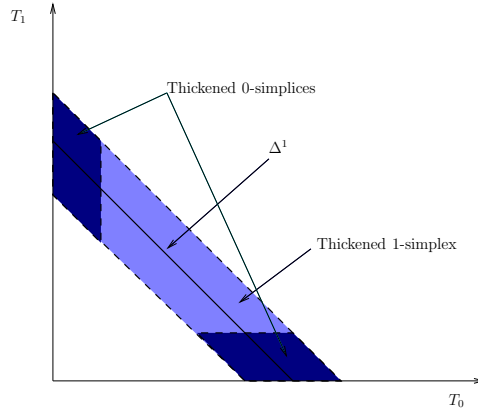


FIGURE 3. Thickening of the simplex Δ^1 .

Notation and definition 3.2. Let $S \subset \mathbf{S}^k \times \mathbf{S}^\ell$ be an open subset defined by the first-order formula $\Phi(\mathbf{X}, \mathbf{Y})$, and let $\pi_{\mathbf{Y}}$ denote the projection along the \mathbf{Y} co-ordinates.

We will denote by $D_{\mathbf{Y},o}^p(\Phi)$ the following first-order formula.

$$(3.3) \quad D_{\mathbf{Y},o}^p(\Phi) \stackrel{\text{def}}{=} \Theta_1(\mathbf{T}) \wedge \Theta_2(\mathbf{X}, \mathbf{Y}^0, \dots, \mathbf{Y}^p, \mathbf{T}) \wedge \Theta_3(U_0, \mathbf{X}, \mathbf{Y}^0, \dots, \mathbf{Y}^p, \mathbf{T})$$

where

$$\begin{aligned} \Theta_1 &\stackrel{\text{def}}{=} \left(\bigwedge_{i=0}^p T_i > 0 \right) \wedge \left(1 - \frac{1}{2(p+1)} < \sum_{i=0}^p T_i < 1 + \frac{1}{2(p+1)} \right), \\ \Theta_2 &\stackrel{\text{def}}{=} \bigwedge_{i=0}^p \left((|\mathbf{Y}^i|^2 < 3/2) \wedge \left((T_i < \frac{1}{2(p+1)}) \vee \Phi_+(\mathbf{X}, \mathbf{Y}^i) \right) \right), \\ \Theta_3 &\stackrel{\text{def}}{=} (U_0^2 + |\mathbf{X}|^2 + \sum_{i=0}^p |\mathbf{Y}^i|^2 + |\mathbf{T}|^2 = 2p + 4) \wedge (U_0 > 0) \end{aligned}$$

and

$$\Phi_+(\mathbf{X}, \mathbf{Y}) \stackrel{\text{def}}{=} (1/2 < |\mathbf{X}|^2 < 3/2) \wedge (1/2 < |\mathbf{Y}|^2 < 3/2) \wedge \Phi(\mathbf{X}/|\mathbf{X}|, \mathbf{Y}/|\mathbf{Y}|).$$

We will denote by $D_{\mathbf{Y},o}^p(S)$ the semi-algebraic set defined by $D_{\mathbf{Y},o}^p(\Phi)$. Notice that $D_{\mathbf{Y},o}^p(S)$ is an open subset of the upper hemisphere of the sphere $\mathbf{S}^N(0, 2p+4)$, where $N = (k+1) + (p+1)(\ell+2)$.

We now prove some important properties of the sets $D_{\mathbf{Y},c}^p(S), D_{\mathbf{Y},o}^p(S)$ defined above as well as of the formulas $D_{\mathbf{Y},c}^p(\Phi), D_{\mathbf{Y},o}^p(\Phi)$ defining them.

Proposition 3.3 (Polynomial time computability). *Suppose there exists a polynomial time real Turing machine M which recognizes the sequence of semi-algebraic sets $(S_n)_{n>0}$ defined by the sequence of first order formulas*

$$(\Phi_n(X_0, \dots, X_{k(n)}, Y_0, \dots, Y_{\ell(n)})_{n>0}, k, \ell = n^{O(1)})$$

where for each $n > 0$, Φ_n defines a closed (respectively open) semi-algebraic subset S_n of $\mathbf{S}^{k(n)} \times \mathbf{S}^{\ell(n)}$. Then there exists a polynomial time real Turing machine M' recognizing the semi-algebraic sets defined by $(D_{\mathbf{Y},c}^p(\Phi_n))_{n>0}$ (respectively $(D_{\mathbf{Y},o}^p(\Phi_n))_{n>0}$).

Proof. Clear from the construction of the formulas $(D_{\mathbf{Y},c}^p(\Phi_n))_{n>0}$ (respectively $(D_{\mathbf{Y},o}^p(\Phi_n))_{n>0}$). \square

We now prove an important topological property of the semi-algebraic sets $D_{\mathbf{Y},c}^p(S), D_{\mathbf{Y},o}^p(S)$ defined above.

Proposition 3.4 (Homotopy equivalence to the join). *Let $S \subset \mathbf{S}^k \times \mathbf{S}^\ell$ be a closed (respectively, open) subset of $\mathbf{S}^k \times \mathbf{S}^\ell$ defined by a first-order formula $\Phi(\mathbf{X}, \mathbf{Y})$, and let $\pi_{\mathbf{Y}}$ denote the projection along the \mathbf{Y} co-ordinates. Then for all $p \geq 0$, $J_{\pi_{\mathbf{Y}}}^p(S)$ is homotopy equivalent to $D_{\mathbf{Y},c}^p(S)$ (respectively, $D_{\mathbf{Y},o}^p(S)$).*

Proof. Suppose S is a closed subset of $\mathbf{S}^k \times \mathbf{S}^\ell$ and let

$$g : D_{\mathbf{Y},c}^p(S) \rightarrow J_{\pi_{\mathbf{Y}}}^p(S)$$

be the map which takes a point $(u, \mathbf{x}, \mathbf{y}^0, \dots, \mathbf{y}^p, \mathbf{t}) \in D_{\mathbf{Y},c}^p(S)$ to the equivalence class represented by the point $((\mathbf{x}, \mathbf{y}^0), \dots, (\mathbf{x}, \mathbf{y}^p), \mathbf{t})$ in $J_{\pi_{\mathbf{Y}}}^p(S)$. From the definition of the spaces $D_{\mathbf{Y},c}^p(S)$ and $J_{\pi_{\mathbf{Y}}}^p(S)$, we have that the inverse image under g of a point represented by $((\mathbf{x}, \mathbf{y}^0), \dots, (\mathbf{x}, \mathbf{y}^p), \mathbf{t})$ in $J_{\pi_{\mathbf{Y}}}^p(S)$ is given by

$$g^{-1}(((\mathbf{x}, \mathbf{y}^0), \dots, (\mathbf{x}, \mathbf{y}^p), \mathbf{t})) = \{(u, \mathbf{x}, \mathbf{z}^0, \dots, \mathbf{z}^p, \mathbf{t}) \mid \text{for each } i, 0 \leq i \leq p, \\ \mathbf{z}^i \in \mathbf{B}^{\ell+1} \text{ and } \mathbf{z}^i = \mathbf{y}^i \text{ if } t_i \neq 0, u^2 + |\mathbf{x}|^2 + \sum_{i=0}^p |\mathbf{z}^i|^2 + |\mathbf{t}|^2 = p + 4, u \geq 0\}.$$

It is easy to see from the above formula that the inverse image under g of each point of $J_{\pi_{\mathbf{Y}}}^p(S)$ is homeomorphic to a product of balls and hence contractible. The proposition now follows from the Vietoris-Begle theorem.

The open case is proved analogously after an infinitesimal retraction reducing it to the closed case. \square

As an immediate corollary we obtain

Corollary 3.5. *Let $S \subset \mathbf{S}^k \times \mathbf{S}^\ell$ be a closed (respectively, open) subset of $\mathbf{S}^k \times \mathbf{S}^\ell$ defined by a first-order formula $\Phi(\mathbf{X}, \mathbf{Y})$, and let $\pi_{\mathbf{Y}}$ denote the projection along the \mathbf{Y} co-ordinates. Then for all $p \geq 0$, $D_{\mathbf{Y},c}^p(S)$ (respectively, $D_{\mathbf{Y},o}^p(S)$) is p -equivalent to $\pi_{\mathbf{Y}}(S)$, and*

$$b_i(D_{\mathbf{Y},c}^p(S)) = b_i(\pi_{\mathbf{Y}}(S))$$

(respectively, $b_i(D_{\mathbf{Y},o}^p(S)) = b_i(\pi_{\mathbf{Y}}(S))$) for $0 \leq i < p$.

Proof. Since S is either an open or closed subset of $\mathbf{S}^k \times \mathbf{S}^\ell$ it is clear that the projection map $\pi_{\mathbf{Y}}$ covers semi-algebraic compacts. Now apply Theorem 2.12. \square

We now show how the formulas $D_{\mathbf{Y},c}^p(\Phi)$ and $D_{\mathbf{Y},o}^p(\Phi)$ can be rewritten when the formula Φ involves quantified blocks of variables.

Lemma 3.6. *Suppose the first-order formula $\Phi(\mathbf{X}, \mathbf{Y})$ is of the form*

$$\Phi \stackrel{\text{def}}{=} (Q_1 \mathbf{Z}^1 \in \mathbf{S}^{k_1})(Q_2 \mathbf{Z}^2 \in \mathbf{S}^{k_2}) \dots (Q_\omega \mathbf{Z}^\omega \in \mathbf{S}^{k_\omega}) \Psi(\mathbf{X}, \mathbf{Y}, \mathbf{Z}^1, \dots, \mathbf{Z}^\omega)$$

with $Q_i \in \{\exists, \forall\}$, and Ψ a quantifier-free first order formula. Let $\pi_{\mathbf{Y}}$ denote the projection along the \mathbf{Y} coordinates. Then, for each $p \geq 0$ the formula

$$D_{\mathbf{Y},*}^p(\Phi)(\mathbf{X}, \mathbf{Y}^0, \dots, \mathbf{Y}^p, \mathbf{T})$$

(where $*$ denotes either c or o) is equivalent to the formula

$$\begin{aligned} \bar{D}_{\mathbf{Y},*}^p(\Phi) &\stackrel{\text{def}}{=} (Q_1 \mathbf{Z}^{1,0} \in \mathbf{S}^{k_1}, \dots, Q_1 \mathbf{Z}^{1,p} \in \mathbf{S}^{k_1}) \\ &\quad (Q_2 \mathbf{Z}^{2,0} \in \mathbf{S}^{k_2}, \dots, Q_2 \mathbf{Z}^{2,p} \in \mathbf{S}^{k_2}) \\ &\quad \vdots \\ &\quad (Q_\omega \mathbf{Z}^{\omega,0} \in \mathbf{S}^{k_\omega}, \dots, Q_\omega \mathbf{Z}^{\omega,p} \in \mathbf{S}^{k_\omega}) \\ &\quad (D_{\mathbf{Y},*}^p(\Psi)(\mathbf{X}, \mathbf{Y}^0, \dots, \mathbf{Y}^p, \mathbf{Z}^{1,0}, \dots, \mathbf{Z}^{\omega,p}, T_0, \dots, T_p)), \end{aligned}$$

where $\mathbf{Y}^i = (Y_0^i, \dots, Y_\ell^i)$ and $\mathbf{Z}^{j,i} = (Z_0^{j,i}, \dots, Z_{k_j}^{j,i})$ for $0 \leq i \leq p, 1 \leq j \leq \omega$, and $\pi_{\mathbf{Y}}$ is the projection along the \mathbf{Y} co-ordinates.

Proof. It follows from the structure of the formula $D_{\mathbf{Y},*}^p(\Phi)(\mathbf{X}, \mathbf{Y}^0, \dots, \mathbf{Y}^p, \mathbf{T})$ that the inner most quantifiers can be pulled outside at the cost of introducing $(p+1)$ copies of the quantified variables. \square

4. PROOF OF THE MAIN THEOREMS

Notation 4.1. Let $\Phi(\mathbf{X})$ be a first-order formula with free variables $\mathbf{X} = (X_1, \dots, X_n)$. We let $\mathcal{R}(\Phi(\mathbf{X}))$ denote the *realization* of the formula Φ ,

$$\mathcal{R}(\Phi(\mathbf{X})) = \{\mathbf{x} \in \mathbb{R}^n \mid \Phi(\mathbf{x})\}.$$

We are now in a position to prove Theorem 1.16. The proof depends on the following key proposition. (Note that we are going to use Proposition 4.2 in the special case when the set of variables Y is empty).

Proposition 4.2. Let $m(n), k_1(n), \dots, k_\omega(n)$ be polynomials, and let $(\Phi_n(\mathbf{X}, \mathbf{Y}))_{n>0}$ be a sequence of formulas

$$\Phi_n(\mathbf{X}, \mathbf{Y}) \stackrel{\text{def}}{=} (Q_1 \mathbf{Z}^1 \in \mathbf{S}^{k_1}) \dots (Q_\omega \mathbf{Z}^\omega \in \mathbf{S}^{k_\omega}) \phi_n(\mathbf{X}, \mathbf{Y}, \mathbf{Z}^1, \dots, \mathbf{Z}^\omega),$$

having free variables $(\mathbf{X}, \mathbf{Y}) = (X_0, \dots, X_{k(n)}, Y_0, \dots, Y_{m(n)})$, with

$$Q_1, \dots, Q_\omega \in \{\exists, \forall\}, Q_i \neq Q_{i+1},$$

where ϕ_n a quantifier-free formula defining a closed (respectively open) semi-algebraic subset of

$$\mathbf{S}^k \times \mathbf{S}^m \times \mathbf{S}^{k_1} \times \dots \times \mathbf{S}^{k_\omega}.$$

Suppose that there exists a real Turing machine M such that for all $n > 0$, M tests membership in $\mathcal{R}(\phi_n(\mathbf{X}, \mathbf{Y}, \mathbf{Z}^1, \dots, \mathbf{Z}^\omega))$ in polynomial time.

Then, there exists $N = N(m)$ a fixed polynomial in n , variables $\mathbf{V} = (V_0, \dots, V_N)$ and a sequence of quantifier-free first order formulas

$$(\Theta_n(\mathbf{X}, \mathbf{V}))_{n>0}$$

such that for each $\mathbf{x} \in \mathbf{S}^{k(n)}$, $\Theta_n(\mathbf{x}, \mathbf{V})$ describes a closed (respectively open) semi-algebraic subset T_n of \mathbf{S}^N . To this sequence $(\Theta_n(\mathbf{X}, \mathbf{V}))_{n>0}$, we can associate polynomial-time computable maps

$$F_n : \mathbb{Z}[T]_{\leq N} \rightarrow \mathbb{Z}[T]_{\leq m}$$

such that the Poincaré polynomials of the fibers over \mathbf{x} verify

$$P_{\mathcal{R}(\Phi_n(\mathbf{x}, \mathbf{Y}))} = F_n(P_{\mathcal{R}(\Theta_n(\mathbf{x}, \mathbf{V}))}).$$

Moreover, there exists a real Turing machine M' testing membership in $\mathcal{R}(\Theta_n(\mathbf{X}, \mathbf{V}))$ in polynomial time.

Proof. The proof is by an induction on ω . We assume that the formula ϕ_n defines a closed semi-algebraic set. The open case can be handled analogously.

If $\omega = 0$ then we let $\Theta_n = \Phi_n$ and $M' = M$, $N = m$, and F_n to be the identity map. Since there are no quantifiers, for each $n \geq 0$ the semi-algebraic set recognized by M and M' are the same and thus the Betti numbers of the sets recognized by M and M' agree.

If $\omega > 0$, we have the following two cases.

(A) Case 1, $Q_1 = \exists$: In this case consider the sequence of formulas $\Phi'_n \stackrel{\text{def}}{=} \bar{D}_{\pi_{\mathbf{Z}^1}, c}^n(\Psi_n)$ (cf. Lemma 3.6), where Ψ_n is the following formula with free variables \mathbf{Y}, \mathbf{Z}^1

$$(4.1) \quad \Psi_n(\mathbf{X}, \mathbf{Y}, \mathbf{Z}^1) \stackrel{\text{def}}{=} (Q_2 \mathbf{Z}^2 \in \mathbf{S}^{k_2}) \cdots (Q_\omega \mathbf{Z}^\omega \in \mathbf{S}^{k_\omega}) \phi_n(\mathbf{X}, \mathbf{Y}, \mathbf{Z}^1, \dots, \mathbf{Z}^\omega).$$

The formula $\bar{D}_{\mathbf{Z}^1, c}^m(\Psi_n)$ is given by

$$\begin{aligned} \bar{D}_{\pi_{\mathbf{Z}^1}, c}^m(\Psi_n) &\stackrel{\text{def}}{=} (Q_2 \mathbf{Z}^{2,0} \in \mathbf{S}^{k_2}, \dots, Q_2 \mathbf{Z}^{2,m} \in \mathbf{S}^{k_2}) \\ &\quad (Q_3 \mathbf{Z}^{3,0} \in \mathbf{S}^{k_3}, \dots, Q_3 \mathbf{Z}^{3,m} \in \mathbf{S}^{k_3}) \\ &\quad \vdots \\ &\quad (Q_\omega \mathbf{Z}^{\omega,0} \in \mathbf{S}^{k_\omega}, \dots, Q_\omega \mathbf{Z}^{\omega,m} \in \mathbf{S}^{k_\omega}) \\ &\quad (D_{\mathbf{Z}^1, c}^m(\phi_n)(\mathbf{X}, \mathbf{Y}, \mathbf{Z}^{1,0}, \dots, \mathbf{Z}^{1,m}, \mathbf{Z}^{2,0}, \dots, \mathbf{Z}^{\omega,m}, T_0, \dots, T_m)). \end{aligned}$$

Note that the quantifier-free inner formula in the above expression,

$$D_{\mathbf{Z}^1, c}^m(\phi_n)(\mathbf{X}, \mathbf{Y}, \mathbf{Z}^{1,0}, \dots, \mathbf{Z}^{1,m}, \mathbf{Z}^{2,0}, \dots, \mathbf{Z}^{\omega,m}, T_0, \dots, T_m)$$

has

$$N = \sum_{j=1}^{\omega} (k_j + 1)(m + 1) + 2(m + 1) = n^{O(1)}$$

free variables, and it is clear from the definition of the formula $D_{\mathbf{Z}^1, c}^m(\phi_n)$ (cf. Eqn. 3.2), that there exists a polynomial time Turing machine (say M_1) to evaluate it since we have a polynomial time Turing machine M for evaluating ϕ_n . Moreover, the formula $\bar{D}_{\pi_{\mathbf{Z}^1}, c}^m(\Psi_n)$ has one less quantifier alternation than the formula Φ_n .

We now apply the proposition inductively to obtain a machine M_2 evaluating $(\Theta_n)_{n>0}$, and a polynomial time computable maps $(G_n)_{n>0}$. By inductive hypothesis we can suppose that for each

$i, 0 \leq i \leq m$ we have for each $\mathbf{x} \in \mathbf{S}^{k(n)}$

$$P_{\mathcal{R}(\bar{D}_{\mathbf{Z}^1, c}^m(\Psi_n(\mathbf{x}, \mathbf{Y}, \mathbf{Z}^1)))} = G_n(P_{\mathcal{R}(\Theta_n(\mathbf{x}, \cdot))}).$$

But by Corollary 3.5, we have that for $0 \leq i \leq m$,

$$b_i(\mathcal{R}(\Phi_n(\mathbf{x}, \mathbf{Y}))) = b_i(\pi_{\mathbf{Z}^1}(\mathcal{R}(\Psi_n(\mathbf{x}, \mathbf{Y}, \mathbf{Z}^1)))) = b_i(\mathcal{R}(\bar{D}_{\mathbf{Z}^1, c}^n(\Psi_n(\mathbf{x}, \mathbf{Y}, \mathbf{Z}^1)))).$$

We set $M' = M_2$ and

$$F_n = G_n.$$

which completes the induction in this case.

- (B) Case 2, $Q_1 = \forall$: In this case consider the sequence of formulas $\Phi'_n \stackrel{\text{def}}{=} \bar{D}_{\mathbf{Z}^1, o}^m(\neg\Psi_n)$ (cf. Lemma 3.6), where Ψ_n is defined as in the previous case (Eqn (4.1)).

We now apply the proposition inductively as above to obtain a machine M_2 evaluating $(\Theta_n)_{n>0}$, and maps $(G_n)_{n>0}$. By inductive hypothesis we can suppose that for each $\mathbf{x} \in \mathbf{S}^n$ we have

$$P_{\mathcal{R}(\bar{D}_{\mathbf{Z}^1, o}^m(\neg\Psi_n(\mathbf{x}, \mathbf{Y}, \mathbf{Z}^1)))} = G_n(P_{\mathcal{R}(\Theta_n(\mathbf{x}, \cdot))}).$$

By Corollary 3.5, we have for $0 \leq i \leq m$,

$$b_i(\mathbf{S}^m \setminus \mathcal{R}(\Phi_n(\mathbf{x}, \mathbf{Y}))) = b_i(\pi_{\mathbf{Z}^1}(\mathcal{R}(\neg\Psi_n(\mathbf{x}, \mathbf{Y}, \mathbf{Z}^1)))) = b_i(\mathcal{R}(\bar{D}_{\mathbf{Z}^1, o}^m(\neg\Psi_n(\mathbf{x}, \mathbf{Y}, \mathbf{Z}^1)))).$$

The set $K = \mathcal{R}(\Phi_n(\mathbf{x}, \mathbf{Y}))$ is a semi-algebraic compact, so by Alexander-Lefschetz duality (Theorem 2.14), we have

$$\begin{aligned} b_0(K) &= 1 + b_{m-1}(\mathbf{S}^m - K) - b_m(\mathbf{S}^m - K), \\ b_i(K) &= b_{m-i-1}(\mathbf{S}^m - K), \quad 1 \leq i \leq m-2, \\ b_{m-1}(K) &= b_0(\mathbf{S}^m - K) - 1 + \max(1 - b_0(\mathbf{S}^m - K), 0), \\ b_m(K) &= 1 - \min(1, b_0(\mathbf{S}^m - K)). \end{aligned}$$

We set $M' = M_2$ and F_n defined by

$$\begin{aligned} F_{n,0}(P) &= 1 + G_{n,m-1}(P) - G_{n,m}(P), \\ F_{n,i}(P) &= G_{n,m-i-1}(P), \quad 1 \leq i \leq m-2, \\ F_{n,m-1}(P) &= G_{n,0}(P) - 1 + \max(1 - G_{n,0}(P), 0), \\ F_{n,m}(P) &= 1 - \min(1, G_{n,0}(P)) \end{aligned}$$

where we denote by $F_{n,i}(P)$ (respectively $G_{n,i}(P)$) the coefficient of T^i in $F_n(P)$ (respectively $G_n(P)$). This completes the induction in this case as well.

□

Proof of Theorem 1.16. Follows immediately from Proposition 4.2 in the special case when the set of variables \mathbf{Y} is empty. In this case the sequence

of formulas $(\Phi_n)_{n>0}$ correspond to a language in the polynomial hierarchy and for each n , $\mathbf{x} = (x_0, \dots, x_{k(n)}) \in S_n \subset \mathbf{S}^{k(n)}$ if and only if

$$F_n(P_{\mathcal{R}(\Theta_n(\mathbf{x}, \cdot))})(0) > 0$$

and this last condition can be checked in polynomial time with advice from the class $\#\mathbf{P}_R^\dagger$. \square

Remark 4.3. It is interesting to observe that in complete analogy with the proof of the classical Toda's theorem the proof of Theorem 1.16 also requires just one call to the oracle at the end.

Proof of Theorem 1.20. Follows from the proof of Proposition 4.2 since the semi-algebraic the formula Θ_n is clearly computable in polynomial time from the given formula Φ_n as long as the number of quantifier alternations ω is bounded by a constant. \square

REFERENCES

1. S. Basu, *On bounding the Betti numbers and computing the Euler characteristic of semi-algebraic sets*, Discrete Comput. Geom. **22** (1999), no. 1, 1–18.
2. ———, *Algorithmic semi-algebraic geometry and topology – recent progress and open problems*, Surveys on Discrete and Computational Geometry: Twenty Years Later, Contemporary Mathematics, vol. 453, American Mathematical Society, 2008, pp. 139–212.
3. S. Basu, R. Pollack, and M.-F. Roy, *Algorithms in real algebraic geometry*, Algorithms and Computation in Mathematics, vol. 10, Springer-Verlag, Berlin, 2006 (second edition). MR 1998147 (2004g:14064)
4. ———, *Computing the first Betti number of a semi-algebraic set*, Found. Comput. Math. **8** (2008), no. 1, 97–136.
5. S. Basu, Richard R. Pollack, and M.-F. Roy, *On the combinatorial and algebraic complexity of quantifier elimination*, J. ACM **43** (1996), no. 6, 1002–1045. MR 98c:03077
6. Saugata Basu, *Computing the first few Betti numbers of semi-algebraic sets in single exponential time*, J. Symbolic Comput. **41** (2006), no. 10, 1125–1154. MR 2262087 (2007k:14120)
7. Saugata Basu, Richard Pollack, and Marie-Françoise Roy, *Computing roadmaps of semi-algebraic sets on a variety*, J. Amer. Math. Soc. **13** (2000), no. 1, 55–82. MR 1685780 (2000h:14048)
8. L. Blum, F. Cucker, M. Shub, and S. Smale, *Complexity and real computation*, Springer-Verlag, New York, 1998, With a foreword by Richard M. Karp. MR 99a:68070
9. L. Blum, M. Shub, and S. Smale, *On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines*, Bull. Amer. Math. Soc. (N.S.) **21** (1989), no. 1, 1–46. MR 90a:68022
10. Peter Bürgisser and Felipe Cucker, *Variations by complexity theorists on three themes of Euler, Bézout, Betti, and Poincaré*, Complexity of computations and proofs (Jan Krajíček, ed.), Quad. Mat., vol. 13, Dept. Math., Seconda Univ. Napoli, Caserta, 2004, pp. 73–151. MR 2131406 (2006c:68053)
11. ———, *Counting complexity classes for numeric computations. II. Algebraic and semialgebraic sets*, J. Complexity **22** (2006), no. 2, 147–191. MR 2200367 (2007b:68059)

12. Peter Bürgisser, Felipe Cucker, and Martin Lotz, *Counting complexity classes for numeric computations. III. Complex projective sets*, Found. Comput. Math. **5** (2005), no. 4, 351–387. MR 2189543 (2006h:68039)
13. J. Canny, *Computing road maps in general semi-algebraic sets*, The Computer Journal **36** (1993), 504–514.
14. Hans Delfs and Manfred Knebusch, *Locally semialgebraic spaces*, Lecture Notes in Mathematics, vol. 1173, Springer-Verlag, Berlin, 1985. MR 819737 (87h:14019)
15. Pierre Deligne, *La conjecture de Weil. I*, Inst. Hautes Études Sci. Publ. Math. (1974), no. 43, 273–307. MR 0340258 (49 #5013)
16. ———, *La conjecture de Weil. II*, Inst. Hautes Études Sci. Publ. Math. (1980), no. 52, 137–252. MR 601520 (83c:14017)
17. B. Dwork, *On the rationality of the zeta function of an algebraic variety*, American Journal of Mathematics **82** (1960), no. 3, 631–648.
18. Mário J. Edmundo and Nicholas J. Peatfield, *o -minimal Čech cohomology*, Q. J. Math. **59** (2008), no. 2, 213–220. MR 2428077
19. A. Gabrielov and N. Vorobjov, *Approximation of definable sets by compact families and upper bounds on homotopy and homology*, preprint at arXiv:math.AG/0710.3028v1, 2007.
20. A. Gabrielov, N. Vorobjov, and T. Zell, *Betti numbers of semialgebraic and sub-Pfaffian sets*, J. London Math. Soc. (2) **69** (2004), no. 1, 27–43. MR 2025325 (2004k:14105)
21. L. Gournay and J. J. Risler, *Construction of roadmaps of semi-algebraic sets*, Appl. Algebra Eng. Commun. Comput. **4** (1993), no. 4, 239–252.
22. D. Yu. Grigor'ev and N. N. Vorobjov, Jr., *Solving systems of polynomial inequalities in subexponential time*, J. Symbolic Comput. **5** (1988), no. 1-2, 37–64. MR 949112 (89h:13001)
23. D. Grigoriev, *Complexity of deciding Tarski algebra*, J. Symbolic Comput. **5** (1988), no. 1-2, 65–108. MR 90b:03054
24. D. Grigoriev and N. Vorobjov, *Counting connected components of a semi-algebraic set in subexponential time*, Comput. Complexity **2** (1992), no. 2, 133–186.
25. Jiří Matoušek, *Using the Borsuk-Ulam theorem*, Universitext, Springer-Verlag, Berlin, 2003, Lectures on topological methods in combinatorics and geometry, Written in cooperation with Anders Björner and Günter M. Ziegler. MR 1988723 (2004i:55001)
26. Klaus Meer, *Counting problems over the reals*, Theoret. Comput. Sci. **242** (2000), no. 1-2, 41–58. MR 1769145 (2002g:68041)
27. C. Papadimitriou, *Computational complexity*, Addison-Wesley, 1994.
28. J. Renegar, *On the computational complexity and geometry of the first-order theory of the reals. I-III.*, J. Symbolic Comput. **13** (1992), no. 3, 255–352.
29. Uwe Schöningh, *Probabilistic complexity classes and lowness*, J. Comput. System Sci. **39** (1989), no. 1, 84–100. MR 1013721 (91b:68041a)
30. Michael Shub and Steve Smale, *On the intractability of Hilbert's Nullstellensatz and an algebraic version of “NP \neq P?”*, Duke Math. J. **81** (1995), no. 1, 47–54 (1996), A celebration of John F. Nash, Jr. MR 1381969 (97h:03067)
31. Stephen Smale, *A Vietoris mapping theorem for homotopy*, Proc. Amer. Math. Soc. **8** (1957), 604–610. MR 0087106 (19,302f)
32. Edwin H. Spanier, *Algebraic topology*, McGraw-Hill Book Co., New York, 1966. MR 0210112 (35 #1007)
33. Larry J. Stockmeyer, *The polynomial-time hierarchy*, Theoret. Comput. Sci. **3** (1976), no. 1, 1–22 (1977). MR 0438810 (55 #11716)
34. A. Tarski, *A decision method for elementary algebra and geometry*, University of California Press, Berkeley and Los Angeles, Calif., 1951, 2nd ed. MR 13,423a
35. Seinosuke Toda, *PP is as hard as the polynomial-time hierarchy*, SIAM J. Comput. **20** (1991), no. 5, 865–877. MR 1115655 (93a:68047)

- 36. L. G. Valiant and V. V. Vazirani, *NP is as easy as detecting unique solutions*, Theoret. Comput. Sci. **47** (1986), no. 1, 85–93. MR 871466 (88i:68021)
- 37. A. Weil, *Number of solutions of equations over finite fields*, Bulletin of the American Mathematical Society **55** (1949), 497–508.

DEPARTMENT OF MATHEMATICS, PURDUE UNIVERSITY, WEST LAFAYETTE, IN 47906,
U.S.A.

E-mail address: `sbasu@math.purdue.edu`

SCHOOL OF MATHEMATICS AND COMPUTING SCIENCES, LENOIR-RHYNE UNIVERSITY,
HICKORY, NC 28603

E-mail address: `thierry.zell@lr.edu`