Implicit Representations via the Polynomial Method

Jean Cardinal

Université libre de Bruxelles (ULB), Belgium

Micha Sharir

School of Computer Science, Tel Aviv University, Israel

Abstract

Semialgebraic graphs are graphs whose vertices are points in $\mathbb{R}^d$, and adjacency between two vertices is determined by the truth value of a semialgebraic predicate of constant complexity. We show how to harness polynomial partitioning methods to construct compact adjacency labeling schemes for families of semialgebraic graphs.

That is, we show that for any family of semialgebraic graphs, given a graph on n vertices in this family, we can assign a label consisting of $O(n^{1-2/(d+1)+\varepsilon})$ bits to each vertex (where $\varepsilon > 0$ can be made arbitrarily small and the constant of proportionality depends on ε and on the complexity of the adjacency-defining predicate), such that adjacency between two vertices can be determined solely from their two labels, without any additional information. We obtain for instance that unit disk graphs and segment intersection graphs have such labelings with labels of $O(n^{1/3+\varepsilon})$ bits. This is in contrast to their natural implicit representation consisting of the coordinates of the disk centers or segment endpoints, which sometimes require exponentially many bits. It also improves on the best known bound of $O(n^{1-1/d}\log n)$ for d-dimensional semialgebraic families due to Alon (*Discrete Comput. Geom.*, 2024), a bound that holds more generally for graphs with shattering functions bounded by a degree-d polynomial. Our labeling scheme is efficient in the sense that not only adjacency between two vertices can be decided in time linear in the size of their labels, but the labels can be computed in subquadratic time on a real RAM from the input points and the semialgebraic adjacency predicate, using recent polynomial partitioning algorithms.

We also give new bounds on the size of adjacency labels for other families of graphs. In particular, we consider semilinear graphs, which are semialgebraic graphs in which the predicate only involves linear polynomials. We show that semilinear graphs have adjacency labels of size $O(\log n)$. We also prove that polygon visibility graphs, which are not semialgebraic in the above sense, have adjacency labels of size $O(\log^3 n)$.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Semialgebraic graphs, Geometric intersection graphs, Visibility graphs, Adjacency labelings, Polynomial partitioning

Digital Object Identifier 10.4230/LIPIcs.WG.2026.10

Related Version *Full Version*: `https://arxiv.org/abs/2602.10922`

Funding *Jean Cardinal*: Work by Jean Cardinal was supported by the Fonds de la Recherche Scientifique-FNRS under Grant n° T003325F.
Micha Sharir: Work by Micha Sharir was partially supported by ISF grant 495/23.

Acknowledgements The first author would like to thank Noga Alon for a discussion on his results in [12], and Lena Yuditsky for sharing observations on semilinear families. The authors are also grateful to Pankaj Agarwal, Sariel Har-Peled, and the WG reviewers for insightful comments and references.

52nd International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2026).
Editors: Jan Goedgebeur and Paweł Rzążewski; Article No. 10; pp. 10:1–10:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

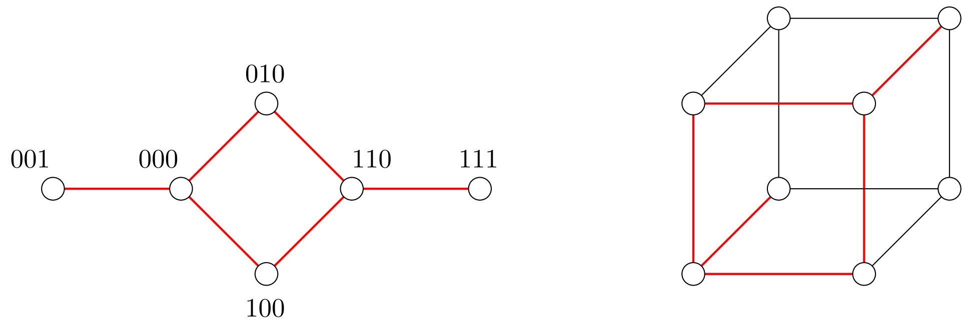

Figure 1 Example of an adjacency labeling of a graph (left). Here two vertices u and v are adjacent, hence $A(\ell(u), \ell(v))$ is true, if and only if the Hamming distance between the two labels $\ell(u)$ and $\ell(v)$ is exactly equal to one. The graph is therefore an induced subgraph of the cube (shown in red, right).

1 Introduction

Finding compact encodings of graphs is a standard topic in computer science and discrete mathematics, see for instance the seminal monograph of Spinrad [70]. In this paper, we explore a special kind of distributed encoding, that associates a binary label, of relatively small size, to each vertex, in such a way that adjacency between two vertices can be deduced solely from their labels, without any additional information about the graph. We focus on graphs defined by semialgebraic predicates, which are commonly encountered in computational geometry and other domains. We first carefully define our objects of study, then summarize our contributions. Our results are stated in Section 1.4.

1.1 Adjacency labeling schemes

We only consider simple, undirected graphs. A family of graphs is said to have a $t(n)$*-bit adjacency labeling scheme*, where n is the number of vertices of the graph, if there exists a computable boolean function $A : \{0,1\}^* \times \{0,1\}^* \mapsto \{\text{true}, \text{false}\}$ such that for every graph (V, E) on n vertices in this family, there exists a map $\ell : V \mapsto \{0,1\}^{t(n)}$ such that $(u, v) \in E$ if and only if $A(\ell(u), \ell(v))$ is true. We may further require that the function A be computable efficiently, that there exist an efficient algorithm to produce the labels given a graph in the family, and that adjacency between two given vertices can be determined efficiently, but the problem makes sense even without these additional requirements. Labeling schemes are also referred to as *implicit representations*, and have been the topic of intensive research since the seminal results of Muller [63] and Kannan, Naor, and Rudich [52]. See Figure 1 for an example.

Note that if a family has a $t(n)$-bit labeling scheme for which the labeling function ℓ is injective, then for every n, there exists a *universal graph* U_n on $2^{t(n)}$ vertices, such that any n-vertex graph in the family is an induced subgraph of U_n [52]. Hence upper bounds on the size of adjacency labeling schemes can be thought of as upper bounds on the size of universal graphs.

Clearly, the family of all graphs has an $O(n)$-bit adjacency labeling scheme. One may, for instance, assign to each vertex the binary word formed by its row in the adjacency matrix of the graph, together with its index (column number) in the matrix. Then, to determine whether (u, v) is an edge, we check whether the v-entry of the row of u is 1. It is actually known that the family of all graphs has an $n/2 + O(1)$ adjacency labeling scheme, which is optimal, up to an additive constant [11, 16]. Results on trees [15], cycles and paths [2], and bounded-degree graphs [1, 3] are also known. It is also known that planar graphs have adjacency labeling schemes of size $(1+o(1))\log_2 n$ [37].

The *speed* of a family of graphs is the function $s(n)$ that maps an integer n to the number of graphs on the vertex set $[n]$ in the family. Observe that if there exists a $t(n)$-bit adjacency labeling scheme for a family of graphs, then $s(n) \leq 2^{nt(n)}$. Therefore, for a family of speed $s(n)$, the function $\frac{1}{n}\log s(n)$ is a lower bound on the minimum size of an adjacency labeling scheme.

A family of graphs is said to be *hereditary* if for any graph in the family, all its induced subgraphs are also in the family. Kannan et al. [52] conjectured that every hereditary family with speed $2^{O(n\log n)}$ has an $O(\log n)$-bit adjacency labeling scheme. Hatami and Hatami [50] refuted the conjecture by showing that for every $\delta > 0$, there is a hereditary family of graphs with speed $2^{O(n\log n)}$ but with no $n^{1/2-\delta}$-bit adjacency labeling scheme.

1.2 Semialgebraic graphs

A simple graph $G = (V, E)$ is a *semialgebraic graph* of description complexity t and dimension d if the vertices in V can be mapped to points in $\mathbb{R}^d$, so that the presence of an edge is defined by a Boolean combination of the sign patterns of t $2d$-variate polynomial functions, of degrees at most t, of the corresponding pair of image points.

More precisely, we have $V \subset \mathbb{R}^d$, and there exist t $2d$-variate polynomials $f_1, \ldots, f_t \in \mathbb{R}[x_1, \ldots, x_d, y_1, \ldots, y_d]$, each of degree at most t, and a Boolean function Φ on $3t$ variables such that for any $u, v \in V$,

$$(u, v) \in E \Leftrightarrow \Phi\left(\{f_i(u, v) < 0, f_i(u, v) = 0, f_i(u, v) \leq 0\}_{i\in[t]}\right).$$

Note that if the graph is undirected, as we generally assume in this work, the predicate has to be symmetric, in the sense that exchanging u and v yields the same truth value. A family of graphs is said to be a *d-dimensional semialgebraic family* if every graph in the family is semialgebraic with the same parameters $d, t, f_1, \ldots, f_t$, and Φ, where d and t are constants (so the dependence on t, the f_i, and Φ is suppressed in this notation).

Semialgebraic families include many well-studied families, such as geometric intersection graphs. Among the simplest examples, one can think of *unit disk graphs*, which are intersection graphs of unit disks in the plane. There every vertex v can be assigned the coordinates of the center of the corresponding disk, and the function Φ simply encodes the condition that two centers are at distance at most two. Many standard problems and tools from graph theory, such as regularity lemmas [42, 44, 65], Erdős-Hajnal properties [14], Ramsey-Turán numbers [45], Zarankiewicz's problem [43], and Ramsey-type results [32] have been studied in a semialgebraic setting.

When the polynomials $f_1, f_2, \ldots, f_t$ in the above definition are all linear, the graphs are said to be *semilinear*. Families of semilinear graphs include many well-studied families of graphs, such as interval graphs, permutation graphs, and bounded-boxicity graphs. Again, several standard problems in graph theory have been studied recently in the semilinear setting [20, 72].

The question of the existence of implicit representations of small size for semialgebraic families has been recently raised by Alon [12, Section 3, paragraph 3]. See Section 1.5 for a discussion of known results.

1.3 Adjacency labeling schemes via biclique decompositions

A *biclique* is a complete bipartite graph. A *biclique decomposition* of a graph is a partition of its edge set into complete bipartite subgraphs. Biclique decompositions are ubiquitous in computational geometry, where they are typically obtained as a byproduct of offline range searching data structures [8, 10, 22, 25, 26, 39, 55, 58]. Biclique decompositions are also useful as compressed representations of graphs on which some algorithms (such as breadth-first search) can be run efficiently [23, 40].

The *size of a biclique decomposition* is the sum of the number of vertices of the bicliques in the decomposition. Note that this is not the same as the number of bicliques of the decomposition, a parameter which has also been studied, and whose behavior, even asymptotic, may be different from the size. The minimum size of a biclique decomposition of a given graph is sometimes referred to as the *representation complexity* of the graph. Indeed, up to logarithmic factors, this size is an upper bound on the number of bits needed to encode the graph; see below. Chung, Erdős, and Spencer [31] proved that every n-vertex graph has a biclique decomposition of size at most $((\log e)/2 + o(1))n^2/\log n$, which is tight up to a factor $1/e$. Erdős and Pyber [38] further showed that every graph has a biclique decomposition such that every vertex belongs to at most $O(n/\log n)$ bicliques of the decomposition. The size of a biclique decomposition also plays a key role in bounds for Zarankiewicz's problem, on the maximum number of edges of bipartite graphs forbidding large bicliques [43]. Indeed, it can easily be shown that if a $K_{t,t}$-free bipartite graph has a biclique cover of size s, then it has at most st edges.

Our key observation is that adjacency labeling schemes can be defined from biclique decompositions as follows: For each vertex v, give the list of bicliques v belongs to, and for each of these bicliques, the side of the bipartition that contains v. The adjacency test simply consists of intersecting the two given lists, and checking that the two vertices are on opposite sides of one of their common bicliques, if any. An example is given in Figure 2.

The length of these labels in bits is $\nu(n) \cdot O(\log n)$, where $\nu(n)$ is an upper bound on the number of bicliques a vertex belongs to. The factor $O(\log n)$ is the number of bits needed to identify a biclique. A single bit per biclique suffices to indicate the side containing the vertex. It is also immediate that by sorting the list of bicliques in increasing order of their labels, adjacency can be tested in time linear in the length of the lists that contain the pair of vertices. The result of Erdős and Pyber [38] yields an $O(n)$ bound on the label size for arbitrary graphs. Note that this labeling scheme needs not be injective. For instance, if the whole graph is a biclique, then only two distinct labels are needed. If we insist on having an injective scheme, we can simply append a $\lceil \log n \rceil$-size identifier of a vertex to its label, so that all labels are distinct.

A related technique, with a constant number of graphs covering each vertex, has been used for adjacency labeling under the name *locally bounded coverings* [18].

1.4 Our results

The main result of this paper is that the recent polynomial partitioning strategies used in biclique decomposition algorithms for semialgebraic graphs can be adapted to ensure that the upper bound $\nu(n)$ on the number of bicliques to which every vertex belongs is strongly sublinear, thereby providing a semialgebraic version of the Erdős-Pyber Theorem [38]. As a consequence, we can give a sublinear bound on adjacency labels for semialgebraic families.

▶ **Theorem 1.** *d-dimensional semialgebraic families have an $O(n^{1-2/(d+1)+\varepsilon})$-bit adjacency labeling scheme, where $\varepsilon > 0$ can be made arbitrarily small and the constant of proportionality depends on ε and on the complexity of the adjacency-defining predicate.*

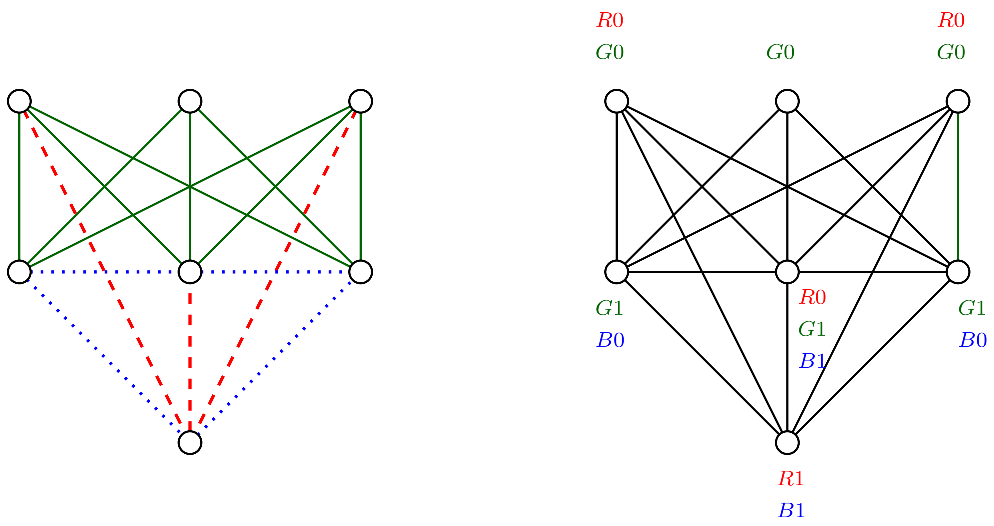

Figure 2 Example of an adjacency labeling of a graph with 16 edges (right) obtained from a biclique decomposition of size 14 (left). The label $\ell(v)$ of a vertex v consists of the list of bicliques v is contained in, from the set $\{R, G, B\}$, together with an additional bit indicating on which side of the biclique the vertex v lies.

This implies the following, among many similar, corollaries:

▶ **Corollary 1.** *Unit disk graphs have an $O(n^{1/3+\varepsilon})$-bit adjacency labeling scheme, for any $\varepsilon > 0$. Disk graphs (for disks with arbitrary radii) have an $O(n^{1/2+\varepsilon})$-bit adjacency labeling scheme, for any $\varepsilon > 0$. Unit ball graphs (intersection graphs for congruent balls) in $\mathbb{R}^d$ have an $O(n^{1-2/(d+1)+\varepsilon})$-bit adjacency labeling scheme, for any $\varepsilon > 0$. Ball graphs (for balls with arbitrary radii) in $\mathbb{R}^d$ have an $O(n^{1-2/(d+2)+\varepsilon})$-bit adjacency labeling scheme, for any $\varepsilon > 0$.*

▶ **Corollary 2.** *Segment intersection graphs have an $O(n^{1/3+\varepsilon})$-bit adjacency labeling scheme, for any $\varepsilon > 0$.*

The second corollary requires justification, since the number of degrees of freedom, also called the *parametric dimension* of segments in the plane is 4. Nevertheless, the predicate that tests whether two segments e and e' intersect one another is a disjunction of conjunctions, each consisting of four polynomial inequalities, each involving at most two parameters of e and at most two parameters of e' (each such inequality determines the side of the line that is spanned by one segment that contains an endpoint of the other segment). In the terminology of [6], we say that the *reduced parametric dimension* of segments, with respect to segment intersection, is only 2. It is easy to see, and will follow from the analysis given later, that the bound in the theorem uses $d = 2$ instead of 4, from which the corollary follows. This observation is general, and calls for a more careful definition of d-dimensional semialgebraic graphs, using the reduced parametric dimension of the objects (see Agarwal *et al.* [6] and Sauermann [66]).

Recall that a graph is semilinear if it is semialgebraic and the t defining polynomials are linear. We prove that semilinear families admit much more concise implicit representations.

▶ **Theorem 2.** *Semilinear families have an $O(\log n)$-bit adjacency labeling scheme.*

Note that many well-studied graph families are semilinear: cographs, interval graphs, permutation graphs, bounded boxicity graphs, circle graphs (intersection graphs of chords of a circle), distance-hereditary graphs (that preserve graph distances under taking connected

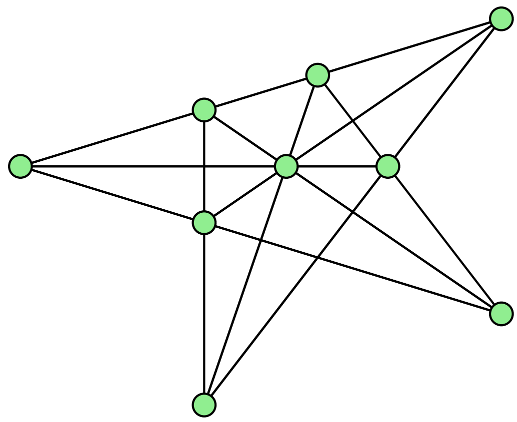

Figure 3 The Perles configuration of 9 points and 9 lines, every realization of which requires at least one irrational coordinate. This provides an example of a semialgebraic family, the bipartite point-line incidence graphs, for which the natural implicit representation by real numbers cannot be turned into an adjacency labeling scheme with labels using a bounded number of bits.

induced subgraphs). While $O(\log n)$-bit adjacency labeling schemes are already known separately for the mentioned classes [70, 33, 19], we obtain this result for all semilinear graphs in an unified way. In a nutshell, the proof of Theorem 2 involves a decomposition of the semilinear graph into a constant number of *comparability graphs* of bounded dimension, as is done in Tomon [72] and Cardinal and Yuditsky [25], each of which is encoded by a simple scheme involving the rank of each vertex in each of the dimensions.

We also consider implicit representations of visibility graphs [35, 46, 64]. *Polygon visibility graphs* are graphs whose vertices are in bijection with the vertices of a simple closed polygon, so that two vertices are adjacent if and only if the line segment between them is contained in the polygon. Note that polygon visibility graphs do not form a semialgebraic family, in the sense defined above, since the adjacency of two vertices depends on the location of the other vertices. In that case, however, a well-known result of Agarwal, Alon, Aronov, and Suri [5] allows us to obtain an efficient labeling scheme.

▶ **Theorem 3.** *Polygon visibility graphs have an $O(\log^3 n)$-bit adjacency labeling scheme.*

We also obtain a slightly better $O(\log^2 n)$ bound for *bichromatic segment intersection graphs*, and for *capped graphs*, a closely related family studied by Davies, Krawczyk, McCarty, and Walczak [34], which contains *terrain visibility graphs*, where the polygon is replaced by an unbounded x-monotone polygonal curve. We refer to the full version of the paper for the complete statements.

1.5 Discussion

Note that semialgebraic graphs are already defined by an implicit representation involving points in $\mathbb{R}^d$. This representation, however, is not easily converted into a binary labeling scheme: There exist semialgebraic graphs, all integer realizations of which require coordinates with doubly exponentially many bits, hence labels of exponentially many bits. This happens even in the special case of simple geometric intersection graphs such as disk and segment intersection graphs, see respectively McDiarmid and Müller [60], and Kratochvíl and Matoušek [56]. This phenomenon can be traced back to early examples of order types requiring coordinates with doubly exponentially many bits, due to Goodman and Pollack [47].

In fact, the situation is even worse: There exist semialgebraic families for which the natural implicit representation by real numbers is not always rational, hence cannot always be (naïvely) encoded by integer numbers. This is the case, for instance, of bipartite point-line incidence graphs, as witnessed by the famous Perles configuration [76] given in Figure 3. These examples all relate to the algebraic universality phenomenon first described by Mnëv [62, 68].

Biclique decompositions of semialgebraic graphs have been studied by Do [36]. She showed that d-dimensional semialgebraic graphs have biclique decompositions of size $O(n^{2d/(d+1)+\varepsilon})$. The same statement, together with construction algorithms, has been given in the context of semialgebraic range searching by Agarwal et al. [6, Appendix A], and in the special case of graphs defined by a single polynomial by Cardinal and Sharir [24]. Note that for a labeling scheme defined as above from a biclique decomposition, the sum of the lengths of the labels over all vertices is equal, up to a logarithmic factor, to the size of the biclique decomposition. The best possible length we can expect to achieve with this method is therefore of the order of $n^{2d/(d+1)+\varepsilon}/n = n^{1-2/(d+1)+\varepsilon}$, for any $\varepsilon > 0$, when every vertex belongs to approximately the same number of bicliques of the decomposition. Our result shows that this is always possible. In fact, we observe that several previously studied methods for constructing biclique decompositions actually produce balanced biclique decompositions, in this sense.

It is known that d-dimensional semialgebraic families have speed $2^{\Theta(n \log n)}$ for any fixed d [66], hence this only provides a logarithmic lower bound on the length of adjacency labels. We leave open the problem of finding any better lower bound on adjacency labeling schemes for semialgebraic families, or alternatively improve on the upper bound stated in Theorem 1.

The existence of sublinear-size labeling schemes for semialgebraic families was known before. Fitch [41] was the first to prove the existence of $O(n^{1-\delta})$-bit adjacency labeling schemes for any semialgebraic family of graphs, for some δ depending on the family. In particular, he gave an upper bound of $O(n^{0.976723})$ for unit disk graphs. Alon [12] proved the following:

▶ **Theorem 4** (Alon [12])**.** *For every $\varepsilon > 0$, there is an integer $d \geq 1$ such that if a hereditary family of graphs has speed at most $2^{(1/4-\varepsilon)n^2}$, then the family has an $O(n^{1-1/d} \log n)$-bit adjacency labeling scheme.*

While this result does not directly address adjacency labelings of semialgebraic graphs, the proof actually yields a similar bound for d-dimensional semialgebraic families. We state the following result, weaker than our Theorem 1, as a consequence of Alon's proof of Theorem 4.

▶ **Theorem 5.** *d-dimensional semialgebraic families have an $O(n^{1-1/d} \log n)$-bit adjacency labeling scheme.*

Proof. Given a collection of binary vectors of length n, the *shatter function* of this family is the function $g(m)$, for $m \in [n]$, that gives the maximum number of distinct projections of the vectors on m coordinates, where the maximum is taken over all choices of m coordinates. Let us now consider the rows of the adjacency matrix of a d-dimensional semialgebraic graph. We claim that these vectors have a polynomially-bounded shatter function.

▶ **Lemma 3.** *The shatter function $g(m)$ of the rows of the adjacency matrix of a d-dimensional semialgebraic graph satisfies $g(m) = O(m^d)$.*

Proof. Let $q_1, \ldots, q_m$ be the points that determine the m coordinates (i.e., columns) of the adjacency matrix. Then, for a point p that determines some row, the signs of the elements of the p-row at columns $q_1, \ldots, q_m$ are determined by the signs of the mt d-variate polynomials (in p) $f_i(p, q_j)$, for $i = 1, \ldots, t$ and $j = 1, \ldots, m$, at the point p. The so-called *sign pattern* of this sequence of polynomials is upper bounded by the number of cells of their arrangement, which, by classical Milnor-Thom type bounds (such as Warren's theorem [73]), is $O((mt)^d) = O(m^d)$, since t is a constant and the polynomials are of constant degree. ◀

We now use the fact that if a family of vectors has a shatter function $g(m) = O(m^d)$, then they can be efficiently encoded in our sense. The following lemma, used by Alon [12], was proved by Welzl [74] and reformulated in those terms by Alon, Moran, and Yehudayoff [13]. For a binary vector v, we refer to a pair of consecutive indices i and $i+1$ such that $v_i \neq v_{i+1}$ as a *switch*.

▶ **Lemma 4.** *Consider a family of binary vectors of length n with shatter function $g(m) \leq c \cdot m^d$ for each m, for some constant $c > 0$ and integer $d \geq 1$. Then there is a fixed permutation of the coordinates of the vectors such that for each permuted vector v, the number of switches in v is at most $O(n^{1-1/d})$.*

The adjacency labeling is obtained by first applying such a permutation to the columns of the adjacency matrix, and then encode each row by the locations of the switches and the sign of the first element. Combining Lemmas 3 and 4 thus yields an upper bound of $O(n^{1-1/d} \log n)$ for d-dimensional semialgebraic families, namely, the sign of the first element, the sequence of switch locations, and the column index of the vertex. This completes the proof of Theorem 5. ◀

In order to improve the $O(n^{1-1/d} \log n)$ bound of Theorem 5 to $O(n^{1-2/(d+1)+\varepsilon})$, as stated in our Theorem 1, we use polynomial partitioning.

1.6 Polynomial partitioning

Geometric partitioning techniques, allowing multidimensional divide-and-conquer algorithms, have been fundamental building blocks in the computational geometry literature since the field's pioneering era. Seminal contributions by Chazelle and Welzl [30], Matoušek [57], and Chazelle [29], among others, have established the standard terminology of partition trees and cuttings. Partition trees were further improved by Chan [27]. We refer to the survey of Agarwal [4] for a thorough account of this line of work through the lens of range searching.

In their paper on the Erdős distinct distances problem, Guth and Katz [49] proposed a novel geometric partitioning scheme that later proved useful for many other problems. Their main statement is the following: For any finite set of points in $\mathbb{R}^d$ and a parameter D, there exists a polynomial $g \in \mathbb{R}[x_1, \ldots, x_d]$ of degree $O(D)$, such that each connected component of $\mathbb{R}^d \setminus Z(g)$, where $Z(g)$ is the zero set of g, contains at most a fraction $1/D^d$ of the points. In a follow-up paper [48], Guth proved that, given a finite set of low-degree k-dimensional varieties in $\mathbb{R}^d$ and a parameter D, there exists a polynomial of degree $O(D)$ such that each component of $\mathbb{R}^d \setminus Z(g)$ intersects at most a fraction $1/D^{d-k}$ of the varieties. These results were later applied successfully in many contexts, including incidences bounds [51, 67, 69], classical theorems in discrete geometry [54], unit distances problems [53, 75], and regularity lemmas [65], to name a few.

Computational aspects of polynomial partitioning, including the delicate recursive handling of situations in which many points of the input set lie on the zero set $Z(g)$ of the partitioning polynomial g (for which the results of [48, 49] provide no useful bound), have been investigated and resolved by Agarwal, Aronov, Ezra, and Zahl [7], and Matoušek and Patáková [59]. This allowed the application of polynomial partitioning techniques to range searching with semialgebraic sets [9], 3SUM-hard geometric problems [17], and algebraic degeneracy testing [24], among others.

Our main result heavily relies on polynomial partitioning, in the extended setups of [7, 59], and provides another perspective on the structure of partition trees produced by the recursive application of the method.

1.7 Plan of the paper

The next section is dedicated to the proof of Theorem 1. The proofs of statements in Theorems 2 and 3 are given in the full version of the paper.

2 Semialgebraic graphs

This section is dedicated to the proof of the following theorem from the introduction, restated here for convenience.

▶ **Theorem 1.** *d-dimensional semialgebraic families have an $O(n^{1-2/(d+1)+\varepsilon})$-bit adjacency labeling scheme, where $\varepsilon > 0$ can be made arbitrarily small and the constant of proportionality depends on ε and on the complexity of the adjacency-defining predicate.*

We first observe that we can restrict our attention to the case where adjacency in the semialgebraic graph is given by the sign of a single polynomial. Indeed, since we assume that the number t of polynomials and the size of the boolean predicate Φ are constant, we can simply encode the vertex by concatenating the t labels corresponding to the t defining polynomials. Then, for given labels λ, λ', we use their sublabels corresponding to the ith polynomial to compute its sign, repeat this for each i, and then apply the function Φ on these signs to determine the adjacency of the two vertices (see for instance Chandoo [28] for related reasonings).

We can also assume, at the cost of a logarithmic factor in the length of labels, that the input is a bipartite semialgebraic graph, in which adjacency is evaluated only between vertices in different parts. We reduce the case of arbitrary graphs to the bipartite case using a simple divide-and-conquer scheme that covers the edges of the input graph by bipartite graphs, so that every vertex is contained in $O(\log n)$ such bipartite graphs.

From now on we will assume that the vertices of the considered semialgebraic graph are points in $\mathbb{R}^d$. Hence the setup is as follows: We are given a bipartite graph $G = (P \cup S, E)$ with $P, S \subset \mathbb{R}^d$ and $E \subseteq P \times S$, and a polynomial $f \in \mathbb{R}[x_1, \ldots, x_d, y_1, \ldots, y_d]$ of constant degree such that $(p, s) \in E$ if and only if $f(p, s) \geq 0$.

We first introduce notations for *ranges* that can be thought of as the duals of the points in P and S. We can map every $s \in S$ to the range

$$s^* := \{x \in \mathbb{R}^d \mid f(x, s) \geq 0\}, \tag{1}$$

so $(p, s) \in E$ if and only if $p \in s^*$. Similarly, with each $p \in P$, we can associate a range

$$p^* := \{y \in \mathbb{R}^d \mid f(p, y) \geq 0\}, \tag{2}$$

and $(p, s) \in E$ if and only if $s \in p^*$.

We then define the dual sets of ranges $P^* := \{p^* \mid p \in P\}$ and $S^* := \{s^* \mid s \in S\}$. The *incidence graph* of a pair (P, R), where P is a set of points and R a set of semialgebraic sets, is the bipartite graph on $P \cup R$ where $(p, r) \in P \times R$ is an edge if and only if $p \in r$. The graph G can therefore be defined as the incidence graph of either the pair (P, S^*) or the pair (S, P^*), where both graphs are defined in $\mathbb{R}^d$.

2.1 Partition trees

Given a pair (P, S^*) consisting of a set P of n points and a set S^* of n semialgebraic sets of constant complexity in $\mathbb{R}^d$, we find a biclique decomposition of the corresponding bipartite incidence graph by constructing a *partition tree* T as follows.

Each node v of T is associated with a pair (P_v, S_v^*), with $P_v \subseteq P$ and $S_v^* \subseteq S^*$. The root node corresponds to the input pair (P, S^*). The child nodes of a node v each contains a subset of P_v, that together form a partition of P_v. Let w be such a child node, with an associated subset $P_w \subset P_v$. Let us consider the set $K_w := \{s^* \in S_v^* \mid P_w \subset s^*\}$ of all ranges in S_v^* that contain all points of P_w. The pair $P_w \times K_w$ is a biclique in the incidence graph of (P_v, S_v^*), which we include in our output. Let us now consider the set $S_w^* := \{s^* \in S_v^* \mid s^* \text{ crosses } P_w\}$, where a range *crosses* a point set if it intersects it but does not contain it. If the number $|P_w|$ of points in P_w is larger than some constant, a partition tree is constructed recursively on the pair (P_w, S_w^*). Otherwise, the remaining $O(|S_w^*|)$ incidences are encoded explicitly. It is clear that all incidences of the incidence graph of (P, S^*) are thus encoded in this way in a one-to-one manner.

The same construction can naturally be defined in the dual setting defined by the pair (S, P^*). Furthermore, the two constructions can be intertwined: Every partitioning step can take place either in the primal or in the dual setting.

For our adjacency labeling purpose, we require that for each vertex of the incidence graph (i.e., point or range), the number of bicliques in which this vertex appears is bounded by an expression that gives, up to a logarithmic factor, the size of the encoding. In the above partition tree, observe that a range appears in a node v if it crosses the corresponding point set P_v, but it participates in a biclique with the point set P_w associated with a child node w of v if it contains P_w. On the other hand, the number of bicliques to which a point belongs is equal to the number of nodes of the partition tree in which it appears. In what follows, since the number of child nodes of a node will be bounded by a constant, the number of bicliques in which a range participates will be proportional to the number of nodes in which it appears. We can therefore restrict our attention to the number of nodes in which each point or range appears.

2.2 Main construction

We construct a partition tree based on the efficient polynomial partitioning techniques described by Matoušek and Patáková [59], and by Agarwal et al. [7]. Our technique is the same as that of Agarwal et al. [6, Appendix A], but the analysis is different, and significantly simpler. In particular, we only focus on balancing the number of nodes in which points and ranges appear in the primal and dual phases, and do not need the full space-query time tradeoff for range searching described in [6]. Another difference is that we restrict the analysis to semialgebraic graphs defined by a single polynomial, which allows us to avoid recursing on the number of conjunctions in the Boolean function Φ written in conjunctive normal form, as done in [6].

The first relevant lemma pertains to the existence of a partitioning polynomial for a given set of points lying in a k-dimensional algebraic variety in $\mathbb{R}^d$, for any $1 \leq k \leq d$. We denote by $Z(g)$ the zero set of a polynomial $g \in \mathbb{R}[x_1, \ldots, x_d]$.

▶ **Lemma 5** (Matoušek and Patáková [59]). *Let V be an algebraic variety of dimension $1 \leq k \leq d$ in $\mathbb{R}^d$ such that all of its irreducible components have dimension k as well, and the degree of every polynomial defining V is at most E. Let $P \subset V$ be a set of n points, and let $D \gg E$ be a parameter.*

There exists a polynomial $g \in \mathbb{R}[x_1, \dots, x_d]$ of degree at most $E^{d^{O(1)}} D$ such that

1. *$V \cap Z(g)$ has dimension at most $k-1$.*
2. *Each connected component of $V \setminus Z(g)$ contains at most n/D^k points of P.*

Assuming D, E, d are constants, the polynomial g, a semialgebraic representation of the connected components of $V \setminus Z(g)$, the points of P lying in each component, as well as the set of points of $P \cap Z(g)$, can be computed in $O(n)$ randomized expected time.

The second lemma generalizes the above result to semialgebraic sets.

▶ **Lemma 6** (Agarwal *et al.* [7]). *Let V be an algebraic variety of dimension $1 \le k \le d$ in $\mathbb{R}^d$, defined by polynomials of degree at most E. Let P^* be a set of n semialgebraic sets in $\mathbb{R}^d$, each of complexity at most some constant b, and let $D \gg E$ be a parameter.*

There exists a polynomial $g \in \mathbb{R}[x_1, \dots, x_d]$ of degree at most $E^{d^{O(1)}} D$ such that

1. *$V \cap Z(g)$ has dimension at most $k-1$.*
2. *Each connected component of $V \setminus Z(g)$ is crossed by (i.e., intersected by but not contained in) at most n/D sets of P^*.*

Assuming D, E, d, and b are constants, the polynomial g, a semialgebraic representation of the connected components of $V \setminus Z(g)$, and the elements of P^* crossing each component can be computed in $O(n)$ randomized expected time.

In what follows, given some partitioning polynomial g, we will refer to the connected components of $V \setminus Z(g)$ as the *cells* of the partition, and to the set $V \cap Z(g)$ as the *special cell*. (The standard cells are all connected, but the special cell does not have to be connected.) The two lemmas above respectively give bounds on the number of points contained in each cell and the number of semialgebraic sets crossing each cell (except for the special cell). A consequence of the well-known Milnor-Thom type series of results [21, 61, 71] is that the number of connected components of $V \setminus Z(g)$, and hence the number of cells, is at most $(cD)^k$ for some constant c depending on d and E. As another consequence of the Milnor-Thom type bounds, we can also give a bound on the number of cells crossed by any semialgebraic set. To state these bounds, we reuse the notations of Lemmas 5 and 6. That is,

▶ **Lemma 7.** *Let V be an algebraic variety of dimension $1 \le k \le d$ in $\mathbb{R}^d$, defined by polynomials of degree at most E, and $g \in \mathbb{R}[x_1, \dots, x_d]$ a polynomial of degree at most $E^{d^{O(1)}} D$, for some parameter $D \gg E$. Then the number of connected components of $V \setminus Z(g)$ crossed by a semialgebraic set of complexity at most b is at most $(cD)^{k-1}$, where the constant c depends on d, E, and b.*

Proof of Theorem 1. We proceed in two phases, corresponding to the primal and dual settings. In the first phase, we construct a partition tree in the primal setting (S, P^*) using Lemma 6. We partition recursively until we reach leaves that contain at most $n^{d/(d+1)}$ ranges from P^*. In the second phase, we construct a partition tree for each such leaf in the dual setting (P, S^*), using Lemma 5, continuing all the way up to leaves with a constant number of points from P. A schematic description of the partition tree thus obtained is given in Figure 4.

In more details, we first consider the pair (S, P^*), where P^* is the set of ranges dual to the points of P. We construct a partition tree by repeatedly applying Lemma 6 with suitably chosen degrees D. Initially, we set V to be the whole space $\mathbb{R}^d$, and put $k = d$. We recurse on the cells defined by the partitioning polynomial g, as well as on the set of points contained in the special cell $V \cap Z(g)$, and a range is retained in the recursive call only if it crosses the cell. In the case of the special cell $V \cap Z(g)$, we do not have any guarantee on

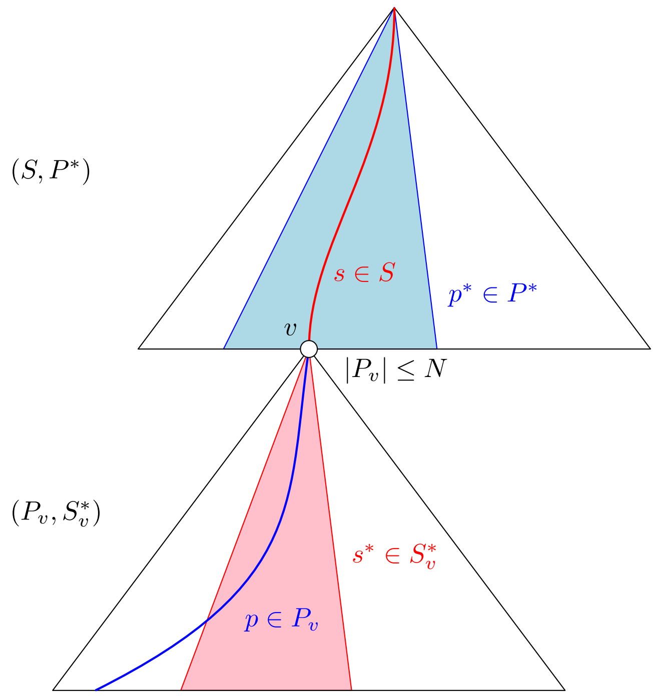

Figure 4 A schematic description of the partition tree used in the proof of Theorem 1. The top tree corresponds to the first phase, in which we repeatedly apply Lemma 6, and at the end of which we are left with leaves each containing at most $N = n^{d/(d+1)}$ elements of P. A dual partition tree is constructed from each leaf v by repeatedly applying Lemma 5. The blue and red subtrees depict the nodes in which, respectively, some element of P (that is, of P^*) and some element of S appear.

the number of points, but Lemma 6 ensures that the dimension k of the ambient variety V decreases. We stop our iterations when the number of semialgebraic sets from P^* remaining in a node is at most $N = n^{d/(d+1)}$.

From Lemma 7, we have that at every cell of this first phase of the partition, the number of child subcells intersected by a set $p^* \in P^*$ is at most $(cD)^{k-1}$, for some constant c (that depends on the much smaller degrees used earlier in the recursion, for higher-degree varieties), which we assume to be ≥ 1, where D is the degree parameter at hand. We can therefore give an upper bound on the number $\nu(n, k)$ of nodes of the partition tree in which a given range $p^* \in P^*$ appears, where n is the size of P^*, and k is the degree of the ambient variety. For $k > 1$, combining Lemma 6 with Lemma 7, we have

$$\nu(n, k) \leq (c_k D_k)^{k-1} \nu(n/D_k, k) + \nu(n, k-1), \tag{3}$$

where the values D_k are the degrees of the partitioning polynomials and each coefficient c_k depends on the earlier degrees, $D_d, \ldots, D_{k+1}$. The degrees D_k are chosen such that $D_{k-1} \gg D_k$. For this choice to make the subsequent analysis work, we require that $D_k^{\varepsilon} \geq 2c_k^{k-1}$. We also have two base cases. First, we have $\nu(N, k) = 1$, for each k, since, by definition, we stop recursing when at most N sets are left. Then we observe that if $k = 1$, we can use balanced binary search trees to get $\nu(n, 1) = O(\log n)$.

The following lemma handles the analysis of the first part of the structure. We stress that in the lemma, the parameter N is fixed, equal to $n^{d/(d+1)}$ for the initial value of n, whereas n varies with the recursion. We refer to the full version of the paper for a proof.

▶ **Lemma 8.** *For each* $1 \le k \le d$ *and each* $n \ge N$*, we have* $\nu(n,k) \le A_k \left(\frac{n}{N}\right)^{k-1+\varepsilon}$*, where* $\varepsilon = \max_{1 \le k \le d} \log_{D_k}(2c_k^{k-1})$*, and* A_k *depends on* ε *and* k*.*

Letting $k = d$ and taking n to be its original value, this solves to:

$$\nu(n,d) \le A_d \left(\frac{n}{N}\right)^{d-1+\varepsilon} = O(n^{1-\frac{2}{d+1}+\varepsilon}).$$

We also observe that every point in S shows up in exactly one subcell at every recursive level, hence the number of nodes of the partition tree containing a given point $s \in S$ is $O(\log n)$.

This describes the first half of our decomposition. We then switch roles and consider the pair (P, S^*), where now S^* is a set of ranges, and P is the original set of points. Each leaf v of the first partition tree corresponds to a set P_v of at most $N = n^{d/(d+1)}$ points of P, and some (possibly large) set S_v^* of ranges from S^*. That is, the first part of the structure guarantees a uniform bound on $|P_v|$ at each of its leaves v, but there is no control on the size of the corresponding set S_v, which could be as large as n. All we know is that $\sum_v |S_v|$, over the leaves v, is n (because the sets S_v are pairwise disjoint).

This calls for another approach, where now we construct another partition tree on each leaf v of the first structure, but now the partition is dictated by the point set P_v and not by the set of ranges, as we did in the first part. Concretely, we repeatedly apply Lemma 5 to the set P_v at a node v, with suitable constants D_k, recursing on both the connected components of $V \setminus Z(g)$ and on the special cell $V \cap Z(g)$, as before. We end the process when the leaves contain only $O(1)$ points of P. Again, we consider the number, now denoted as $\nu^*(n,k)$, of nodes of the partition tree in which a given range $s^* \in S_v^*$ appears, where n is the number of elements in P_v. Combining Lemmas 5 and 7, we obtain the following recursion for $k > 1$, which bears some resemblance to (3):

$$\nu^*(n,k) \le (c_k D_k)^{k-1} \nu^*(n/D_k^k, k) + \nu^*(n, k-1),$$

where the values D_k are the degrees of the partitioning polynomials, chosen such that $D_{k-1} \gg D_k$, as before, and each value c_k depends, as before, on the preceding degrees D_j, for $j > k$. We also have the base cases $\nu^*(1,k) = 1$ and $\nu^*(n,1) = O(\log n)$.

The analog of Lemma 8 in the new context is:

▶ **Lemma 9.** *For each* $1 \le k \le d$ *and each* n*, we have* $\nu^*(n,k) \le B_k n^{1-1/k+\varepsilon}$*, where* $\varepsilon = \max_{1 \le k \le d} \left(\frac{1}{k} \log_{D_k}\left(2^{1/k} c_k^{1-1/k}\right)\right)$*, and* B_k *depends on* ε *and* k*.*

The proof is given in the full version of the paper. Letting $k = d$ and replacing n by $N = n^{d/(d+1)}$, which is the initial value of $|P_v|$ for the second part of the structure, we obtain that every range $s^* \in S_v^*$ appears in at most $O(n^{1-\frac{2}{d+1}+\frac{d\varepsilon}{d+1}})$ nodes.

We may now account for the contributions of the two phases together. In the partition tree of the first phase, every point of S appears in $O(\log n)$ nodes along a path, and every range of P^* appears in $O(n^{1-\frac{2}{d+1}+\varepsilon})$ nodes, by Lemma 8. Then at every leaf v of this partition tree, we attach a dual partition tree on the set of points P_v and the set of ranges S_v^*. In that tree, every point of P_v appears in at most $O(\log n)$ nodes along a path, and every range of S_v^* appears in at most $O(n^{1-\frac{2}{d+1}+\frac{d\varepsilon}{d+1}})$ nodes, by Lemma 9. (We refer to Figure 4 for a schematic description.) So overall, every element of S appears in $O(\log n + n^{1-\frac{2}{d+1}+\varepsilon}) = O(n^{1-\frac{2}{d+1}+\varepsilon})$ nodes, and every element of P appears in $O(n^{1-\frac{2}{d+1}+\frac{d\varepsilon}{d+1}} \log n) = O(n^{1-\frac{2}{d+1}+\varepsilon})$ nodes, with a suitable adjustment of the coefficients, as desired. This concludes the proof of Theorem 1. ◀

References

1 Mikkel Abrahamsen, Stephen Alstrup, Jacob Holm, Mathias Bæk Tejs Knudsen, and Morten Stöckel. Near-optimal induced universal graphs for bounded degree graphs. In *44th International Colloquium on Automata, Languages, and Programming*, volume 80, pages Art. No. 128, 14, 2017. doi:10.4230/LIPIcs.ICALP.2017.128.

2 Mikkel Abrahamsen, Stephen Alstrup, Jacob Holm, Mathias Bæk Tejs Knudsen, and Morten Stöckel. Near-optimal induced universal graphs for cycles and paths. *Discrete Appl. Math.*, 282:1–13, 2020. doi:10.1016/j.dam.2019.10.030.

3 David Adjiashvili and Noy Rotbart. Labeling schemes for bounded degree graphs. In *Automata, languages, and programming. Part II*, volume 8573 of *Lecture Notes in Comput. Sci.*, pages 375–386. Springer, Heidelberg, 2014. doi:10.1007/978-3-662-43951-7_32.

4 Pankaj K. Agarwal. Simplex range searching and its variants: a review. In *A Journey Through Discrete Mathematics. A tribute to Jiří Matoušek*, pages 1–30. Springer, Cham, 2017. doi:10.1007/978-3-319-44479-6_1.

5 Pankaj K. Agarwal, Noga Alon, Boris Aronov, and Subhash Suri. Can visibility graphs be represented compactly? *Discrete Comput. Geom.*, 12(3):347–365, 1994. doi:10.1007/BF02574385.

6 Pankaj K. Agarwal, Boris Aronov, Esther Ezra, Matthew J. Katz, and Micha Sharir. Intersection queries for flat semi-algebraic objects in three dimensions and related problems. *ACM Transactions on Algorithms*, 21(3):25:1–25:59, 2025. doi:10.1145/3721290.

7 Pankaj K. Agarwal, Boris Aronov, Esther Ezra, and Joshua Zahl. Efficient algorithm for generalized polynomial partitioning and its applications. *SIAM J. Comput.*, 50(2):760–787, 2021. doi:10.1137/19M1268550.

8 Pankaj K. Agarwal, Esther Ezra, and Micha Sharir. Semi-algebraic off-line range searching and biclique partitions in the plane. In *40th International Symposium on Computational Geometry*, pages 4:1–4:15, 2024. doi:10.4230/LIPIcs.SOCG.2024.4.

9 Pankaj K. Agarwal, Jiří Matoušek, and Micha Sharir. On range searching with semialgebraic sets. II. *SIAM J. Comput.*, 42(6):2039–2062, 2013. doi:10.1137/120890855.

10 Pankaj K. Agarwal and Kasturi R. Varadarajan. Efficient algorithms for approximating polygonal chains. *Discrete Comput. Geom.*, 23(2):273–291, 2000. doi:10.1007/PL00009500.

11 Noga Alon. Asymptotically optimal induced universal graphs. *Geom. Funct. Anal.*, 27(1):1–32, 2017. doi:10.1007/s00039-017-0396-9.

12 Noga Alon. Implicit representation of sparse hereditary families. *Discrete Comput. Geom.*, 72(2):476–482, 2024. doi:10.1007/s00454-023-00521-0.

13 Noga Alon, Shay Moran, and Amir Yehudayoff. Sign rank and the Vapnik-Chervonenkis dimension. *Mat. Sb.*, 208(12):4–41, 2017. doi:10.4213/sm8780.

14 Noga Alon, János Pach, Rom Pinchasi, Radoš Radoičić, and Micha Sharir. Crossing patterns of semi-algebraic sets. *J. Combin. Theory Ser. A*, 111(2):310–326, 2005. doi:10.1016/j.jcta.2004.12.008.

15 Stephen Alstrup, Søren Dahlgaard, and Mathias Bæk Tejs Knudsen. Optimal induced universal graphs and adjacency labeling for trees. *J. ACM*, 64(4):Art. 27, 22, 2017. doi:10.1145/3088513.

16 Stephen Alstrup, Haim Kaplan, Mikkel Thorup, and Uri Zwick. Adjacency labeling schemes and induced-universal graphs. *SIAM J. Discrete Math.*, 33(1):116–137, 2019. doi:10.1137/16M1105967.

17 Boris Aronov, Mark de Berg, Jean Cardinal, Esther Ezra, John Iacono, and Micha Sharir. Subquadratic algorithms for some 3SUM-hard geometric problems in the algebraic decision-tree model. *Comput. Geom.*, 109:Paper No. 101945, 21, 2023. doi:10.1016/j.comgeo.2022.101945.

18 Aistis Atminas, Andrew Collins, Vadim Lozin, and Victor Zamaraev. Implicit representations and factorial properties of graphs. *Discrete Math.*, 338(2):164–179, 2015. doi:10.1016/j.disc.2014.09.008.

19 Avah Banerjee. An adjacency labeling scheme based on a decomposition of trees into caterpillars. In *Combinatorial algorithms*, volume 13270 of *Lecture Notes in Comput. Sci.*, pages 114–127. Springer, Cham, 2022. doi:10.1007/978-3-031-06678-8_9.

20 Abdul Basit, Artem Chernikov, Sergei Starchenko, Terence Tao, and Chieu-Minh Tran. Zarankiewicz's problem for semilinear hypergraphs. *Forum Math. Sigma*, 9:Paper No. e59, 23, 2021. doi:10.1017/fms.2021.52.

21 Saugata Basu, Richard Pollack, and Marie-Françoise Roy. On the number of cells defined by a family of polynomials on a variety. *Mathematika*, 43(1):120–126, 1996. doi:10.1112/S0025579300011621.

22 Peter Braß and Christian Knauer. On counting point-hyperplane incidences. *Computational Geometry*, 25(1-2):13–20, 2003. doi:10.1016/S0925-7721(02)00127-X.

23 Sergio Cabello, Siu-Wing Cheng, Otfried Cheong, and Christian Knauer. Geometric matching and bottleneck problems. In *40th International Symposium on Computational Geometry*, volume 293, pages 31:1–31:15, 2024. doi:10.4230/LIPIcs.SOCG.2024.31.

24 Jean Cardinal and Micha Sharir. Improved algebraic degeneracy testing. *Discrete Comput. Geom.*, 74(1):177–195, 2025. doi:10.1007/s00454-024-00673-7.

25 Jean Cardinal and Yelena Yuditsky. Compact Representation of Semilinear and Terrain-Like Graphs. In *33rd Annual European Symposium on Algorithms (ESA 2025)*, volume 351 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 67:1–67:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.ESA.2025.67.

26 Timothy M. Chan. Dynamic subgraph connectivity with geometric applications. *SIAM Journal on Computing*, 36(3):681–694, 2006. doi:10.1137/S009753970343912X.

27 Timothy M. Chan. Optimal partition trees. *Discrete Comput. Geom.*, 47(4):661–690, 2012. doi:10.1007/s00454-012-9410-z.

28 Maurice Chandoo. Logical labeling schemes. *Discrete Math.*, 346(10):Paper No. 113565, 21, 2023. doi:10.1016/j.disc.2023.113565.

29 Bernard Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(2):145–158, 1993. doi:10.1007/BF02189314.

30 Bernard Chazelle and Emo Welzl. Quasi-optimal range searching in spaces of finite VC-dimension. *Discrete Comput. Geom.*, 4(5):467–489, 1989. doi:10.1007/BF02187743.

31 Fan R. K. Chung, Paul Erdős, and Joel Spencer. On the decomposition of graphs into complete bipartite subgraphs. In *Studies in pure mathematics: To the Memory of Paul Turán*, pages 95–101. Birkhäuser, Basel, 1983. doi:10.1007/978-3-0348-5438-2_10.

32 David Conlon, Jacob Fox, János Pach, Benny Sudakov, and Andrew Suk. Ramsey-type results for semi-algebraic relations. *Trans. Amer. Math. Soc.*, 366(9):5043–5065, 2014. doi:10.1090/S0002-9947-2014-06179-5.

33 B. Courcelle and R. Vanicat. Query efficient implementation of graphs of bounded clique-width. *Discrete Appl. Math.*, 131(1):129–150, 2003. The Second International Colloquium "Journées de l'Informatique Messine" (Metz, 2000). doi:10.1016/S0166-218X(02)00421-3.

34 James Davies, Tomasz Krawczyk, Rose McCarty, and Bartosz Walczak. Coloring polygon visibility graphs and their generalizations. *J. Combin. Theory Ser. B*, 161:268–300, 2023. doi:10.1016/j.jctb.2023.02.010.

35 Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. Visibility graphs. In *Computational Geometry. Algorithms and Applications, Third Edition*, chapter 15. Springer, 2008. doi:10.1007/978-3-540-77974-2.

36 Thao Do. Representation complexities of semialgebraic graphs. *SIAM J. Discrete Math.*, 33(4):1864–1877, 2019. doi:10.1137/18M1221606.

37 Vida Dujmović, Louis Esperet, Cyril Gavoille, Gwenaël Joret, Piotr Micek, and Pat Morin. Adjacency labelling for planar graphs (and beyond). *J. ACM*, 68(6):Art. 42, 33, 2021. doi:10.1145/3477542.

38 Paul Erdős and László Pyber. Covering a graph by complete bipartite graphs. *Discrete Math.*, 170(1-3):249–251, 1997. doi:10.1016/S0012-365X(96)00124-0.

39 Jeff Erickson. New lower bounds for Hopcroft's problem. *Discrete Comput. Geom.*, 16(4):389–418, 1996. doi:10.1007/BF02712875.

40 Tomás Feder and Rajeev Motwani. Clique partitions, graph compression and speeding-up algorithms. *J. Comput. System Sci.*, 51(2):261–272, 1995. doi:10.1006/jcss.1995.1065.

41 Matthew Fitch. Implicit representation conjecture for semi-algebraic graphs. *Discrete Appl. Math.*, 259:53–62, 2019. doi:10.1016/j.dam.2018.11.030.

42 Jacob Fox, Mikhail Gromov, Vincent Lafforgue, Assaf Naor, and János Pach. Overlap properties of geometric expanders. *J. Reine Angew. Math.*, 671:49–83, 2012. doi:10.1515/crelle.2011.157.

43 Jacob Fox, János Pach, Adam Sheffer, Andrew Suk, and Joshua Zahl. A semi-algebraic version of Zarankiewicz's problem. *J. Eur. Math. Soc. (JEMS)*, 19(6):1785–1810, 2017. doi:10.4171/JEMS/705.

44 Jacob Fox, János Pach, and Andrew Suk. A polynomial regularity lemma for semialgebraic hypergraphs and its applications in geometry and property testing. *SIAM J. Comput.*, 45(6):2199–2223, 2016. doi:10.1137/15M1007355.

45 Jacob Fox, János Pach, and Andrew Suk. Ramsey-Turán numbers for semi-algebraic graphs. *Electron. J. Combin.*, 25(4):Paper No. 4.61, 5, 2018. doi:10.37236/7988.

46 Subir K. Ghosh. *Visibility Algorithms in the Plane.* Cambridge University Press, 2007. doi:10.1017/CBO9780511543340.

47 Jacob E. Goodman, Richard Pollack, and Bernd Sturmfels. Coordinate representation of order types requires exponential storage. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, pages 405–410, 1989. doi:10.1145/73007.73046.

48 Larry Guth. Polynomial partitioning for a set of varieties. *Math. Proc. Cambridge Philos. Soc.*, 159(3):459–469, 2015. doi:10.1017/S0305004115000468.

49 Larry Guth and Nets Hawk Katz. On the Erdős distinct distances problem in the plane. *Ann. of Math. (2)*, 181(1):155–190, 2015. doi:10.4007/annals.2015.181.1.2.

50 Hamed Hatami and Pooya Hatami. The implicit graph conjecture is false. In *63rd Annual IEEE Symposium on Foundations of Computer Science*, pages 1134–1137. IEEE, 2022. doi:10.1109/FOCS54457.2022.00109.

51 Saarik Kalia, Micha Sharir, Noam Solomon, and Ben Yang. Generalizations of the Szemerédi-Trotter theorem. *Discrete Comput. Geom.*, 55(3):571–593, 2016. doi:10.1007/s00454-016-9759-5.

52 Sampath Kannan, Moni Naor, and Steven Rudich. Implicit representation of graphs. *SIAM J. Discrete Math.*, 5(4):596–603, 1992. doi:10.1137/0405049.

53 Haim Kaplan, Jiří Matoušek, Zuzana Safernová, and Micha Sharir. Unit distances in three dimensions. *Combin. Probab. Comput.*, 21(4):597–610, 2012. doi:10.1017/S0963548312000144.

54 Haim Kaplan, Jiří Matoušek, and Micha Sharir. Simple proofs of classical theorems in discrete geometry via the Guth-Katz polynomial partitioning technique. *Discrete Comput. Geom.*, 48(3):499–517, 2012. doi:10.1007/s00454-012-9443-3.

55 Matthew J. Katz and Micha Sharir. An expander-based approach to geometric optimization. *SIAM J. Comput.*, 26(5):1384–1408, 1997. doi:10.1137/S0097539794268649.

56 Jan Kratochvíl and Jiří Matoušek. Intersection graphs of segments. *J. Combin. Theory Ser. B*, 62(2):289–315, 1994. doi:10.1006/jctb.1994.1071.

57 Jiří Matoušek. Efficient partition trees. *Discrete Comput. Geom.*, 8(3):315–334, 1992. doi:10.1007/BF02293051.

58 Jiří Matoušek. Range searching with efficient hierarchical cuttings. *Discrete Comput. Geom.*, 10(2):157–182, 1993. doi:10.1007/BF02573972.

59 Jiří Matoušek and Zuzana Patáková. Multilevel polynomial partitions and simplified range searching. *Discrete Comput. Geom.*, 54(1):22–41, 2015. doi:10.1007/s00454-015-9701-2.

60 Colin McDiarmid and Tobias Müller. Integer realizations of disk and segment graphs. *J. Combin. Theory Ser. B*, 103(1):114–143, 2013. doi:10.1016/j.jctb.2012.09.004.

61 John Milnor. On the Betti numbers of real varieties. *Proc. Amer. Math. Soc.*, 15:275–280, 1964. doi:10.2307/2034050.

62 Nicolai E. Mnëv. The universality theorems on the classification problem of configuration varieties and convex polytopes varieties. In *Topology and Geometry—Rohlin Seminar*, volume 1346 of *Lecture Notes in Math.*, pages 527–543. Springer, Berlin, 1988. doi:10.1007/BFb0082792.

63 John H. Muller. *Local structure in graph classes.* PhD thesis, Georgia Institute of Technology, USA, 1988.

64 Joseph O'Rourke. Visibility. In *Handbook of Discrete and Computational Geometry, 3rd Edition*, chapter 33. Chapman and Hall/CRC, 2017. doi:10.1201/9781315119601.

65 Natan Rubin. An efficient regularity lemma for semi-algebraic hypergraphs. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3611–3636. SIAM, Philadelphia, PA, 2025. doi:10.1137/1.9781611978322.120.

66 Lisa Sauermann. On the speed of algebraically defined graph classes. *Adv. Math.*, 380:Paper No. 107593, 55, 2021. doi:10.1016/j.aim.2021.107593.

67 Micha Sharir and Noam Solomon. Incidences between points and lines in $\mathbb{R}^4$. *Discrete Comput. Geom.*, 57(3):702–756, 2017. doi:10.1007/s00454-016-9822-2.

68 Peter W. Shor. Stretchability of pseudolines is NP-hard. In *Applied Geometry and Discrete Mathematics*, volume 4 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 531–554. Amer. Math. Soc., Providence, RI, 1991. doi:10.1090/dimacs/004/41.

69 József Solymosi and Terence Tao. An incidence theorem in higher dimensions. *Discrete Comput. Geom.*, 48(2):255–280, 2012. doi:10.1007/s00454-012-9420-x.

70 Jeremy P. Spinrad. *Efficient graph representations*, volume 19 of *Fields Institute Monographs*. American Mathematical Society, Providence, RI, 2003. doi:10.1090/fim/019.

71 René Thom. Sur l'homologie des variétés algébriques réelles. In *Differential and Combinatorial Topology (A Symposium in Honor of Marston Morse)*, pages 255–265. Princeton Univ. Press, Princeton, NJ, 1965. doi:10.1515/9781400874842-016.

72 István Tomon. Ramsey properties of semilinear graphs. *Israel J. Math.*, 254(1):113–139, 2023. doi:10.1007/s11856-022-2390-7.

73 Hugh E. Warren. Lower bounds for approximation by nonlinear manifolds. *Trans. Amer. Math. Soc.*, 133:167–178, 1968. doi:10.2307/1994937.

74 Emo Welzl. Partition trees for triangle counting and other range searching problems. In *Proceedings of the 4th Annual Symposium on Computational Geometry*, pages 23–33, 1988. doi:10.1145/73393.73397.

75 Joshua Zahl. An improved bound on the number of point-surface incidences in three dimensions. *Contrib. Discrete Math.*, 8(1):100–121, 2013. doi:10.11575/CDM.V8I1.62120.

76 Günter M. Ziegler. Nonrational configurations, polytopes, and surfaces. *Math. Intelligencer*, 30(3):36–42, 2008. doi:10.1007/BF02985377.