Bidimensionality of Geometric Intersection Graphs[1]

Alexander Grigoriev[2]   Athanassios Koutsonas[3]   Dimitrios M. Thilikos[3,4,5]

Abstract

Let $\mathcal{B}$ be a finite collection of geometric (not necessarily convex) bodies in the plane. Clearly, this class of geometric objects naturally generalizes the class of disks, lines, ellipsoids, and even convex polygons. We consider geometric intersection graphs $G_{\mathcal{B}}$ where each body of the collection $\mathcal{B}$ is represented by a vertex, and two vertices of $G_{\mathcal{B}}$ are adjacent if the intersection of the corresponding bodies is non-empty. For such graph classes and under natural restrictions on their maximum degree or subgraph exclusion, we prove that the relation between their treewidth and the maximum size of a grid minor is linear. These combinatorial results vastly extend the applicability of all the meta-algorithmic results of the bidimensionality theory to geometrically defined graph classes.

Keywords: Geometric intersection graphs, Grid exlusion theorem, Bidimensionality

1   Introduction

Parameterized complexity treats problems as subsets of $\Sigma^* \times \mathbb{N}$, for some alphabet Σ. An instance of a parameterized problem is a pair (I, k) where I is the main part of the problem description and k is a, typically small, parameter. An FPT *algorithm* for a parameterized problem Π is one that runs in $f(k) \cdot n^{O(1)}$ time. A central issue in parameterized complexity is to find which parameterized problems admit FPT algorithms and, when this is the case, to reduce as much as possible the contribution of the function $f(\cdot)$, i.e., their *parametric dependance*. FPT algorithms where $f(k) = 2^{o(k)}$ are

called *sub-exponential parameterized algorithms*. It is known that such an algorithm where $f(k) = 2^{o(\sqrt{k})}$ is unlikely to exist for several problems on graphs, even when restricted to sparse graph classes such as planar graphs [2]. Therefore, a parametric dependance $f(k) = 2^{O(\sqrt{k})}$ is the best we may expect and this is what we may aim for.

A kernelization algorithm for a parameterized problem Π is one that, in polynomial time, can replace any instance (I, k) with a new equivalent one whose size depends exclusively on the parameter k. If such an algorithm exists and the size of the new instance is linear in k, then we say that Π *admits a linear kernel*. While the existence of an FTP algorithm implies the existence of a kernel it is a challenge to find for which problems such a kernel can be polynomial [1].

Bidimensionality theory. This theory was initially introduced in [4] as a general framework for designing parameterized algorithms with sub-exponential parametric dependance. Moreover, it also provided meta-algorithmic results in approximation algorithms [5, 8] and kernelization [10] (for a survey on bidimensionality, see [3]). To present the consequences and the motivation of our results let us first give some brief description of the meta-algorithmic consequences of Bidimensionality theory. For this, we need first some definitions.

A *graph invariant* is a function $\mathbf{p}$ mapping graphs to non-negative integers. The parameterized problem associated to an invariant $\mathbf{p}$ has as input a pair (G, k) where G is a graph and k is a non-negative integer, and asks whether $\mathbf{p}(G) \leq k$ (or, alternatively whether $\mathbf{p}(G) \geq k$). Let $\mathcal{G}$ be the set of all graphs. The parameterized problem corresponding to $\mathbf{p}$ is denoted by $\Pi_{\mathbf{p}} \subseteq \mathcal{G} \times \mathbb{N}$ and is defined as $\Pi_{\mathbf{p}} = \{(G, k) \mid \mathbf{p}(G) \leq k\}$ or, alternatively, as $\Pi_{\mathbf{p}} = \{(G, k) \mid \mathbf{p}(G) \geq k\}$. We also define the graph invariant $\mathbf{bg}$ such that given a graph G,

$$\mathbf{bg}(G) = \max\{k \mid G \text{ contains the } (k \times k)\text{-grid as a minor}\}.$$

Definition 1. *Given a* graph invariant $\mathbf{p}$ *we say that* $\Pi_{\mathbf{p}}$ *is* minor-bidimensional *if the following conditions hold:*

- $\mathbf{p}$ *is closed under taking of subgraphs, i.e., for every* $G \in \mathcal{G}$, *if* H *is a minor of* G, *then* $\mathbf{p}(H) \leq \mathbf{p}(G)$.

- *If* L_k *is the* $(k \times k)$-*grid, then* $\mathbf{p}(L_k) = \Omega(k^2)$.

The main consequences of bidimensionality theory for minor closed invariants are summarized by the following: Suppose that $\mathbf{p}$ is a graph invariant such that $\Pi_{\mathbf{p}}$ is a minor bidimensional problem. Let also $\mathcal{G}$ be a graph class such that it satisfies the following property:

$$\forall_{G \in \mathcal{G}} \; \mathbf{tw}(G) \;\; = \;\; O(\mathbf{bg}(G)) \tag{1}$$

and let $\Pi_{\mathbf{p}}^{\mathcal{G}}$ be the restriction of $\Pi_{\mathbf{p}}$ to the graphs in $\mathcal{G}$, i.e. the problem occurring if we alter all YES-instance of $\Pi_{\mathbf{p}}^{\mathcal{G}}$ whose graph is not in $\mathcal{G}$ to NO-instances. Then the following hold

1. if $\mathbf{p}(G)$ can be computed in $2^{O(\mathbf{tw}(G))} \cdot n^{O(1)}$ steps, then $\Pi_{\mathbf{p}}^{\mathcal{G}}$ can be solved by a sub-exponential parameterized algorithm that runs in $2^{O(\sqrt{k})} \cdot n^{O(1)}$ steps.

2. if $\mathbf{p}$ satisfies some separability property (see [8, 5, 10] for the precise definition) and $\Pi_{\mathbf{p}} = \{(G, k) \mid \exists S \subseteq V(G) : |S| \geq k \text{ and } (G, S) \models \psi\}$ where ψ is a sentence in Counting Monadic Second Order logic, then Π_{p} admits a linear kernel, i.e. there exists a polynomial algorithm reducing each instance (G, k) of Π_{p} to an equivalence instance (G', k') where $|V(G')| = O(k)$ and $k' \leq k$.

3. If $\mathbf{p}$ satisfies some separability property and is reducible (in the sense this is defined in [8]), then there is an EPTAS for computing $\mathbf{p}(G)$ on the graphs in $\mathcal{G}$.

According the the current state of the art all above meta-algorithmic results hold when $\mathcal{G}$ excludes graphs with some fixed graph H as a minor. This is due to the combinatorial result of Demaine and Hajiaghayi in [6], who proved (1) for every graph G excluding some fixed graph H as a minor. While such graphs are of a topological nature it remained an interesting question whether the applicability of the above theory can be extended for geometrically (rather than topologically) restricted graphs classes.

Our results. Clearly, any extension of the applicability of bidimensionality theory on some class $\mathcal{G}$ requires a proof that it satisfies property (1). Recently, a first step to extend meta-algorithmic results for graph classes that are not topologically restricted was done in [9], where the bidimensionality condition was used to derive sub-exponential algorithms for H-free unit-disk intersection graphs and H-free map graphs, where a graph class is H-free if none of its graphs contains H as a subgraph. However, no meta-algorithmic results were known so far for more generic classes of geometric intersection graphs, like e.g. intersection graphs of polygonal objects in the plane. In this paper we vastly extend the combinatorial results of [9] to more general families of geometric intersection graphs. In particular, we prove that property (1) holds for several classes of geometric intersection graphs and open a new direction of the applicability of bidimensionality theory. In particular our results are the following.

1. Let $\mathcal{B}$ be a set of (not necessarily straight) lines in the plane such that for each $C_1, C_2 \in \mathcal{B}$ with $C_1 \neq C_2$, the set $C_1 \cap C_2$ is a finite set of points and at most two lines intersect in the same point. Assume also that each line is intersected at most ξ times. Then $\mathbf{tw}(G_{\mathcal{B}}) = O(\xi \cdot \mathbf{bg}(G_{\mathcal{B}}))$.

2. Let $\mathcal{B}$ be a set of ρ-convex bodies (bodies where any two of their points can be joined by a polysegment of at most $\rho - 1$ bends that is entirely inside the body) such that for each $B_1, B_2 \in \mathcal{B}$ with $B_1 \neq B_2$, the set $B_1 \cap B_2$ has a non-empty interior. Let $G_{\mathcal{B}}$ be the intersection graph of $\mathcal{B}$ and let Δ be the maximum degree of $G_{\mathcal{B}}$. Then $\mathbf{tw}(G_{\mathcal{B}}) = O(\rho^2 \Delta^3 \cdot \mathbf{bg}(G_{\mathcal{B}}))$.

3. Let H be a graph on h vertices, and let $\mathcal{B}$ be a collection of convex bodies in the plane such that for each $B_1, B_2 \in \mathcal{B}$ with $B_1 \neq B_2$, the set $B_1 \cap B_2$ has a non-empty interior. If the intersection graph $G_{\mathcal{B}}$ of $\mathcal{B}$ is α-fat and does not contain H as a subgraph, then $\mathbf{tw}(G_{\mathcal{B}}) = O(\alpha^6 h^3 \cdot \mathbf{bg}(G_{\mathcal{B}}))$. (Given a real number α, we call the intersection graph of a collection of convex bodies α-fat if the ratio between the maximum and the minimum radius of a circle where all bodies of the collection can be circumscribed, and inscribed respectively, is upper bounded by α.)

Notice that the case of H-subgraph free unit-disk intersection graphs treated in [9] is just a very special case of the fourth result (unit-disk graphs are 1-convex and 1-fat).

The paper is organized as follows: In Section 2, we give some basic definitions and results. In Section 3 we prove the main technical results that are used in Section 4 for the derivation of its implications in a variety of geometric graph classes. Section 5 discusses extensions and conclusions of this work. Al proofs have been moved to the appendix except from those of Lemmata 6, 7 and Theorem 1.

2   Definitions and preliminaries

All graphs in this paper are undirected and may have loops or multiple edges. If a graph has no multiple edges or loops we call it *simple*. Given a graph G, we denote by $V(G)$ its vertex set and by $E(G)$ its edge set. Let x be a vertex or an edge of a graph G and likewise for y; their distance in G, denoted by $\mathbf{dist}_G(x, y)$ is the smallest length of a path in G that contains them both. We call *part of a path* any sequence of adjacent edges in a given path. For any set of vertices $S \subseteq V(G)$, we denote by $G[S]$ the subgraph of G induced by the vertices from S.

Graph embeddings. We use the term *graph embedding* to denote a drawing of a graph G in the plane, where each vertex is associated to a distinct point of the plane and each edge to a simple open Jordan curve, such that its endpoints are the two points of the plane associated with the endvertices of this edge. To simplify the presentation, when not necessary, we do not distinguish between a vertex of G and the point in the plane representing the vertex; likewise for an edge of G. Roughly speaking, we often do not distinguish between G and its embedding. Two edges of an embedding of a graph in the plane *cross*, if they share a non-vertex point of the plane. We use the term *plane graph* for an embedding of a graph without crossings. A graph is *planar* if it admits a plane embedding.

Geometric bodies, lines and polysegments. We call a set of points in the plane a *2-dimensional geometric body*, or simply a *2-dimensional body*, if it is homeomorphic to the closed disk $\{(x, y) \mid x^2 + y^2 \leq 1\}$. Also a *line* is a subset of the plane that is homeomorphic to the interval $[0, 1]$. A *polysegment C* is a line that is the union of

a sequence of straight lines $\overline{p_1 p_2}, \overline{p_2 p_3}, \cdots, \overline{p_{k-1} p_k}$ in the plane, where p_1 and p_k are the endpoints of C. We say that a polysegment C *contains* a point p_i and *joins* the endpoints p_1, p_k, and we refer to the rest points $p_2, p_3, \cdots, p_{k-1}$ as *bend points* of C. The length of a polysegment is defined as equal to the number of straight lines it contains (i.e. one more than the number of its bend points). Throughout the paper we assume that a polysegment is not self-crossing.

Minors and distance minors. Given two graphs H and G, we write $H \preccurlyeq G$ and call H a *minor* of G, if H can be obtained from a subgraph of G by edge contractions (the *contraction* of an edge $e = \{x, y\}$ in a graph G is the operation of replacing x and y by a new vertex x_e that is made adjacent with all the neighbors of x and y in G that are different from x and y). Moreover, we say that H is a *contraction* of G, if H can be obtained from G by contracting edges.

Let G be a simple graph. We denote as G^ℓ the graph obtained from G by adding a loop on each of its vertices. We also say that a subset F of $E(G^\ell)$ is *solid*, if for every $v_1, v_2 \in \bigcup_{e \in F} e$ there is a walk in G^ℓ from v_1 to v_2 consisting of edges in F and where each second edge is a loop. We define the relation $\preccurlyeq_\phi$ between two graphs as follows.

Let H and G be simple graphs. Then we write $H \preccurlyeq_\phi G$, if there is a function $\phi : E(G^\ell) \to V(H) \cup E(H) \cup \{\star\}$, such that

1. for every vertex $v \in V(H)$, $\phi^{-1}(v)$ is a non-empty solid set,

2. for every two distinct vertices $v_1, v_2 \in V(H)$, an edge in $\phi^{-1}(v_1)$ does not share a common endpoint with an edge in $\phi^{-1}(v_2)$.

3. for every edge $e = \{v_1, v_2\} \in E(H)$ and every edge e' in $\phi^{-1}(e)$, e' is not a loop and shares its one endpoint with an edge in $\phi^{-1}(v_1)$ and the other with an edge in $\phi^{-1}(v_2)$.

4. for every $e \in E(H)$, $|\phi^{-1}(e)| = 1$.

The following lemma reveals the equivalence between the relation defined previously and the minor relation (for the proofs see Appendix).

Lemma 1. *If G and H are graphs, then $H \preccurlyeq_\phi G$ if and only if H is a minor of G.*

Proof. Let $H \preccurlyeq_\phi G$. First notice, that by the definition of the function ϕ any loop of G^ℓ will be either discarded or mapped to a vertex of H. Conditions (1) and (2) guarantee that any two vertices x, y of H are corresponding to vertex-disjoint connected subgraphs G_x, G_y of G. Moreover by (3), if $xy \in E(H)$ there is an edge in G joining G_x and G_y. Hence, H can be obtained from a subgraph of G by contracting the edges of the subgraphs G_z ($z \in V(H)$) and thus is a minor of G.

Let now H be a minor of G. This means that there exists a subgraph of G consisting of disjoint trees $\{T_v \mid v \in V(H)\}$ plus a set of edges $E = E(H)$, such that contracting all the edges of the trees yields the graph H. Then, we choose ϕ as a function that maps the edges of a tree T_x as well as the loops on the vertices of T_x to x, the edges of E to $E(H)$ and all other edges of G^ℓ to $\star$. $\qquad \square$

Given the existence of a function ϕ as in the definition above, we say H *is a ϕ-generated minor of G*. Moreover, H *is a distance minor of G* if H is a ϕ-generated minor of G and the following additional condition holds:

5. for every $e_1, e_2 \in E(G) \setminus \phi^{-1}(\star)$, $\mathbf{dist}_H(\phi(e_1), \phi(e_2)) \leq \mathbf{dist}_G(e_1, e_2)$.

Contractions and c-contractions. If the definition of the relation $\preccurlyeq_\phi$ is modified by omitting condition (4) and demanding that $\phi^{-1}(\star) = \emptyset$, then we deal with the contraction relation and we say that H *is a ϕ-generated contraction of G*. (Note, that condition (4) is not a requirement of the equivalence to the minor relation – see also the proof of Lemma 1.) Let c be a non negative integer. We say that H is a *c-contraction of G* if H is a ϕ-generated contraction of G and for all $v \in V(H)$, $G[\phi^{-1}(v)]$ is a graph of at most c edges.

In this paper we we use the alternative, more complicated, definitions of minors and contractions as they are necessary for the proofs or our results.

Tree-decompositions and treewidth. A *tree-decomposition* of a graph G, is a pair $(T, \mathcal{X})$, where T is a tree and $\mathcal{X} = \{X_t : t \in V(T)\}$ is a family of subsets of $V(G)$, called *bags*, such that the following three properties are satisfied:

(1) $\bigcup_{t \in V(T)} X_t = V(G)$,

(2) for every edge $e \in E(G)$ there exists $t \in V(T)$ such that X_t contains both ends of e, and

(3) $\forall v \in V(G)$, the set $T_v = \{t \in V(T) \mid v \in X_t\}$ induces a tree in T.

The *width* of a tree-decomposition is the cardinality of the maximum size bag minus 1 and the *treewidth* of a graph G is the minimum width of a tree-decomposition of G. We denote the treewidth of G by $\mathbf{tw}(G)$. We say that a graph H is a *partial triangulation* of a plane graph G if G is a spanning subgraph of H and H is plane. The following result follows from [11].

Proposition 1. *Let r be an integer. Then, any planar graph with treewidth at least $4.5 \cdot r$ contains a partial triangulation of the $(r \times r)$-grid as a contraction.*

Lemma 2. *Let G be a planar graph and k an integer. If $\mathbf{tw}(G) \geq 18 \cdot k$ then G contains a $(k \times k)$-grid as a distance minor.*

Proof. By Proposition 1, the graph G contains a partial triangulation P of a $(4k \times 4k)$ grid as a contraction. We claim that the $(k \times k)$-grid L_k is a ϕ-generated distance minor of P, where the function $\phi : E(P^\ell) \to V(L_k) \cup E(L_k) \cup \{\star\}$ is defined as follows. For any vertex $v_{i,j}$ of L_k, with $i, j \in \{1, \ldots, k\}$, the set $\phi^{-1}(v)$ contains the three loops on the vertices of P with coordinates $(k+2i, k+2j)$, $(k+2i-1, k+2j)$, $(k+2i, k+2j-1)$ and the two edges joining these three vertices. The set $\phi^{-1}(E(L_k))$ contains all possible candidates from the edges of the underlying grid of P, while anything else is mapped to $\star$.

One can easily verify that L_k is indeed the $(k \times k)$-grid, and that the function ϕ satisfies the first four conditions. It remains to show that condition (5) also holds. It suffices to show it for the case of the distance of vertices and the rest follows. Consider, for this purpose, any two vertices v_1, v_2 in L_k and let their distance be $\mathbf{dist}_{L_k}(v_1, v_2) = \rho \leq 2k - 2$. Since any vertex in P that corresponds to a vertex of L_k has a distance of at least k from a boundary vertex, the shortest path in P containing $\phi^{-1}(v_1)$ and $\phi^{-1}(v_2)$ has length at least the half of the length of the shortest path in the underlying grid, i.e. $\frac{1}{2} \cdot 2\rho \geq \rho$. $\qquad\square$

3   Bidimensionality of line intersection graphs

Let $\mathcal{B} = \{B_1, \ldots, B_k\}$ be a collection of lines in the plane. The *intersection graph* $G_\mathcal{B}$ of $\mathcal{B}$, is a graph whose vertex set is $\mathcal{B}$, and that has an edge $\{B_i, B_j\}$ (for $i \neq j$) if and only if B_i and B_j *touch*, namely $B_i \cap B_j \neq \emptyset$.

The following theorem states our main technical result.

Theorem 1. *Let $\mathcal{B}$ be a set of lines in the plane such that for each $C_1, C_2 \in \mathcal{B}$ with $C_1 \neq C_2$, the set $C_1 \cap C_2$ is a finite set of points and at most two lines intersect in the same point. Let also $G_\mathcal{B}$ be the intersection graph of $\mathcal{B}$ and let $\xi = \max_{C \in \mathcal{B}} |C \cap \bigcup_{C' \in \mathcal{B} \setminus C} C'|$. Then $\mathbf{tw}(G_\mathcal{B}) = O(\xi \cdot \mathbf{bg}(G_\mathcal{B}))$.*

To prove Theorem 1 we will need a series of lemmata.

Lemma 3. *Let G be a graph and let H be a c-contraction of G. Then $\mathbf{tw}(G) \leq (c + 1) \cdot (\mathbf{tw}(H) + 1) - 1$.*

Proof. By definition, since H is a c-contraction of G, there is a mapping between each vertex of H and a connected set of at most c edges in G, so that by contracting these edgesets we obtain H from G. The endpoints of these edges form disjoint connected sets in G, implying a partition of the vertices of G into connected sets $\{V_x \mid x \in V(H)\}$, where $|V_x| \leq c + 1$ for any vertex $x \in V(H)$.

Consider now a tree decomposition $(T, \mathcal{W})$ of H. We claim that the pair $(T, \mathcal{W}')$, where $W_t' := \bigcup_{x \in W_t} V_x$ for $t \in T$ is a tree decomposition of G. Clearly all vertices of G are included in some bag, since all vertices of H did. Every edge of G with both endpoints in the same part of the partition is in a bag, as each of these vertex sets is placed as a whole in the same bag. If e is an edge of G with endpoints in different parts of the partition, say V_x and V_y, then this implies that $xy \in E(H)$. Thus, there is a node t of T for which $x, y \in W_t$ and therefore $e \in W_t'$. Moreover, the continuity property remains unaffected, since for any vertex $x \in V(H)$ all vertices in V_x induce the same subtree in T that x did. $\qquad\square$

Lemma 4. *Let G be a graph and let $V_1, \ldots, V_r$ be a partition of the vertices of G, such that for each $i \in \{1, \ldots, r\}$, $G[V_i]$ is a connected graph, and for each $i \in \{1, \ldots, r-1\}$ there exist an edge of G with one endpoint in V_i and one endpoint in V_{i+1}. Let also*

$s \in V_1$ and $t \in V_r$. Then G has a path from s to t with a part P of length at least $\beta - \alpha + 2$, where $1 \le \alpha < \beta \le r$, so that P does not contain any edge in $G[V_i]$ for $i \in \{1, \ldots, \alpha - 1\} \cup \{\beta + 1, \ldots, r\}$.

Proof. For each $i \in \{1, \ldots, r-1\}$ let $e_i = t_i s_{i+1}$ be an edge of G with one endpoint t_i in V_i and the other s_{i+1} in V_{i+1} and set $s_1 = s$ and $t_r = t$. For any $i \in \{1, \ldots, r\}$, since $G[V_i]$ is a connected graph, there is a path P_i from s_i to t_i that lies entirely in $G[V_i]$ (possibly the trivial path of no edges). Then the path $P_1 e_1 \ldots P_{r-1} e_{r-1} P_r$ from s to t, with its part $e_a P_{a+1} \ldots e_{\beta-1} P_\beta e_\beta$ satisfies the requirements of the Lemma. $\qquad\square\qquad\square$

Lemma 5. *Let A, B, and C be graphs such that B is a ψ_1-generated contraction of A and C is a ψ_2-generated minor of A for some functions $\psi_1 : E(A^\ell) \to V(B) \cup E(B)$ and $\psi_2 : E(A^\ell) \to V(C) \cup E(C) \cup \{\star\}$. If*

$$\forall_{e \in E(C)} \qquad |\psi_2^{-1}(e) \cap \psi_1^{-1}(E(B))| = 1 \qquad (2)$$

then C is also a minor of B.

Proof. We will define a function $\phi : E(B^\ell) \to V(C) \cup E(C) \cup \{\star\}$, which in turn guaranties that C is a minor of B. First of all notice that by (2) there is a subset $F_E \subseteq E(B)$ and a bijection η between F_E and $E(C)$. In fact, for any edge e of C it holds that $\eta^{-1}(e) = \psi_1(\psi_2^{-1}(e))$. We set $\phi|_{F_E} = \eta$ and we observe that ϕ is obliged to map any edge in $F' = E(B^\ell) \setminus F_E$ either to a vertex of C or to $\star$.

Let now v be a vertex of $V(C)$. Recall that $\psi_2^{-1}(v)$ is a non-empty solid set of edges in A^ℓ, and thus it induces a connected subgraph, say A_v, in A. Notice also that each edge e_v of C incident to v has a unique pre-image in A, which has exactly one endpoint in A_v. Furthermore, the graph $B_v = B[\psi_1(\psi_2^{-1}(v))]$ is isomorphic to the graph taken if we contract in A_v all its edges that belong in $\psi_1^{-1}(V(B))$. As the contraction of edges does not harm the connectivity of an edgeset, it follows that $\psi_1(\psi_2^{-1}(v))$ is a connected set of edges in B and that again B_v is connected and each edge $\psi_1(\psi_2^{-1}(e_v))$ has an endpoint in B_v. We set $\phi(f) = v$, if $f \in E(B^\ell)$ is an edge of B_v or a loop on a vertex of B_v. Let $F_V \subseteq E(B^\ell)$ be the union of all edgesets $\phi^{-1}(v)$, for each $v \in V(C)$ and observe that $F_E \cap F_V = \emptyset$. Finally, we set $\phi(E(B^\ell) \setminus (F_E \cup F_V)) = \star$. It remains to prove that the four conditions of the minor definition are satisfied.

The first and fourth conditions follow straightforwardly from the definition of ϕ. For the second, assume that v_1, v_2 are two distinct vertices of $V(C)$. We set $N_A = \psi_2^{-1}(v_1) \cup \psi_2^{-1}(v_2)$ and $N_B = \psi_1(\psi_2^{-1}(v_1)) \cup \psi_1(\psi_2^{-1}(v_2))$. Notice that $B[N_B]$ is isomorphic to the graph taken from $A[N_A]$ after we contract all its edges that belong in $\psi_1^{-1}(V(B))$. As $A[N_A]$ is a disconnected graph, the same holds for $B[N_B]$. Therefore N_B is disconnected.

For the third, let $e = (v_1, v_2) \in E(C)$. Clearly $\phi^{-1}(e)$ is not a loop since it belongs in F_E. Moreover, in A, one endpoint of $e' = \psi_2^{-1}(e)$ is in $A[\psi_2^{-1}(v_1)]$ and the other is in $A[\psi_2^{-1}(v_2)]$. Thus, $M_A = \psi_2^{-1}(v_1) \cup \{e'\} \cup \psi_2^{-1}(v_2)$ is a connected set of A while $M_A \setminus \{e'\}$ is not. We set $M_B = \psi_1(\psi_2^{-1}(v_1)) \cup \{e''\} \cup \psi_1(\psi_2^{-1}(v_2))$, where $e'' = \phi^{-1}(e)$

and we observe that $B[M_B]$ is isomorphic to the graph taken from $A[M_A]$ after we contract all its edges that belong in $\psi_1^{-1}(V(B))$ in a way that, in this isomorphism, e' corresponds to the edge e''. This means that M_B is a connected set of B while $M_B \setminus \{e''\}$ is not. We conclude that the one endpoint of e'' belongs in $\phi^{-1}(v_1)$ and the other belongs in $\phi^{-1}(v_2)$, as required. $\qquad\square$

Lemma 6. *Let G be a connected graph and let H be a c-contraction of G. If G contains a $(k \times k)$-grid as a distance minor, then H contains a (k', k')-grid as a minor, where $k' = \lfloor \frac{k-1}{2(c+1)} \rfloor + 1$.*

Proof. We assume that c is an odd number and equivalently prove the lemma for $k' = \lfloor \frac{k-1}{2c} \rfloor + 1$.

Let H be a σ-generated contraction of G for some $\sigma : E(G^\ell) \to V(H) \cup E(H)$ such that $G[\sigma^{-1}(v)]$ is a graph of at most c edges for all $v \in V(H)$. Suppose also that G contains a $(k \times k)$-grid L_k as a distance minor via a function $\phi : E(G^\ell) \to V(L_k) \cup E(L_k) \cup \{\star\}$.

We assume that $V(L_k) = \{1, \ldots, k\}^2$ where each (i, j) corresponds to its grid coordinates. Our target is to prove that the $(k' \times k')$ grid $L_{k'}$ is a minor of H. We define $\alpha : \{1, \ldots, k'\} \to \{1, \ldots, k\}$ such that $\alpha(i) = 2(i-1)c + 1$. Notice that this definition is possible as $2(k'-1)c + 1 \leq k$. For each $(i, j) \in \{1, \ldots, k'\}^2$, we define a *horizontal* and a *vertical* set of vertices in $V(L_k)$,

$$U_{i,j}^{\text{hor}} = \bigcup_{r \in \{\alpha(i)+(c+1)/2,\ldots,\alpha(i+1)-(c+1)/2\}} (r, \alpha(j)), \tag{3}$$

$$U_{i,j}^{\text{ver}} = \bigcup_{r \in \{\alpha(j)+(c+1)/2,\ldots,\alpha(j+1)-(c+1)/2\}} (\alpha(i), r) \tag{4}$$

and let $\mathcal{U}$ be the collection of all sets $U_{i,j}^{\text{hor}}$ or $U_{i,j}^{\text{ver}}$ defined in (3) and (4). For every horizontal (resp. vertical) $U \in \mathcal{U}$, we denote by $\mathcal{E}(U) \subseteq E(L_k)$ the set containing all horizontal (resp. vertical) edges of L_k with an endpoint in U. We will prove the following claim:

(∗) *Let U_1 and U_2 be two different sets of $\mathcal{U}$ and let e_1, e_2 be two edges of G such that $\phi(e_i) \in \mathcal{E}(U_i) \cup U_i$, for $i = \{1, 2\}$. Then, there are no disjoint paths of length at most c from the endpoints of e_1 to the endpoints of e_2 in G.*

Since L_k is a distance minor of G, it suffices to show that there is no cycle in L_k, that contains $\phi(e_1)$ and $\phi(e_2)$ together with two paths between them of length at most c. Let us suppose that such a cycle exists. Notice that by the definition of $\mathcal{U}$, if two vertices x, y of $V(L_k)$ belong to two different sets of $\mathcal{U}$, then $\mathbf{dist}_{L_k}(x, y) \geq c + 1$. This implies that $\phi(e_i)$ must be an edge $v_i u_i$ of L_k with only one endpoint, say u_i, in U_i, for $i = \{1, 2\}$, or else we are done. Likewise, it holds that $\mathbf{dist}_{L_k}(u_1, u_2) \geq c + 1$ and hence one path of length at most c of the cycle must be from v_1 to u_2, the other from v_2 to u_1. It follows, that the edges $v_1 u_1$ and $v_2 u_2$ cannot be both vertical nor horizontal,
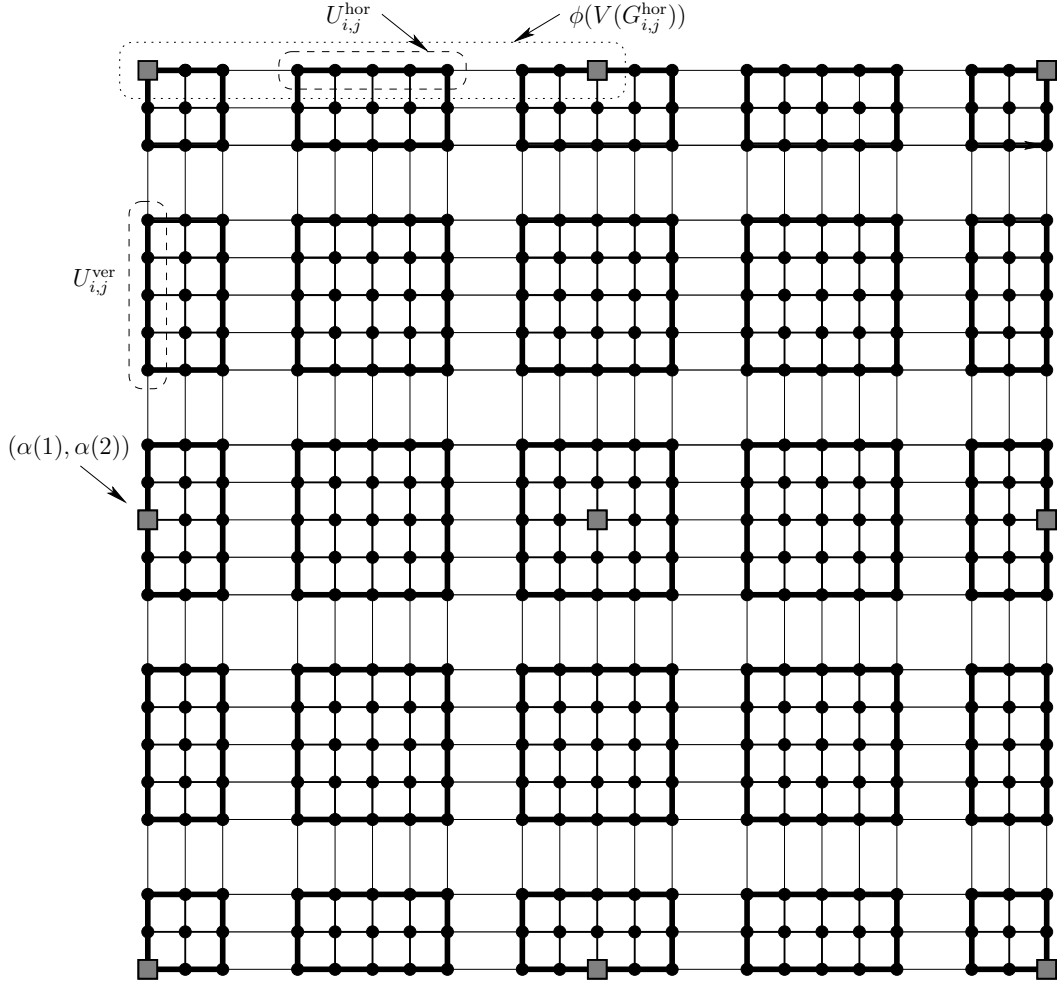
Figure 1: An example of the proof of Lemma 6 for $c = 5$, $k = 21$, and $k' = 3$.

and all vertices of the two disjoint paths lie inside the square part of the grid these two edges define. This contradicts the planarity of the grid, which completes the proof of the claim.

For every $i, j \in \{1, \ldots, k'\}^2$ we choose arbitrarily a vertex $v_{i,j}$ from the graph $G[\phi^{-1}(\alpha(i), \alpha(j))]$. This selection creates a collection of $k' \times k'$ vertices of G.

For each pair $\{(i, j), (i+1, j)\}$ where $(i, j) \in \{1, \ldots, k'-1\} \times \{1, \ldots, k'\}$, we observe that the graph

$$G_{i,j}^{\mathrm{hor}} = G[\bigcup_{i' \in \{\alpha(i), \ldots, \alpha(i+1)\}} \phi^{-1}(i', \alpha(j))]$$

is connected, and for every $i' \in \{\alpha(i), \ldots, \alpha(i+1)\}$ the sets $\phi^{-1}(i', \alpha(j))$ form a partition of $V(G_{i,j}^{\mathrm{hor}})$ and there is an edge of G between $\phi^{-1}(i', \alpha(j))$ and $\phi^{-1}(i'+1, \alpha(j))$. From Lemma 4, $G_{i,j}^{\mathrm{hor}}$ contains a path $P_{i,j}^{\mathrm{hor}}$ from $v_{i,j}$ to $v_{i+1,j}$ with a part of length at least $c + 1$ in $\phi^{-1}(\mathcal{E}(U_{i,j}^{\mathrm{hor}})) \cup \phi^{-1}(U_{i,j}^{\mathrm{hor}})$. Clearly, one of the edges in this part of the path, say $e_{i,j}^{\mathrm{hor}}$, is an edge of $\sigma^{-1}(E(H))$. We denote by $\overrightarrow{P}_{i,j}$ (resp. $\overleftarrow{P}_{i+1,y}$) the part of $P_{i,j}^{\mathrm{hor}}$

starting from $v_{i,y}$ (resp. $v_{i+1,y}$) and containing only one endpoint of $e_{i,j}^{\text{hor}}$.

Working in the same way as before but following the "vertical" instead of "horizontal" direction, for each pair $\{(i,j),(i,j+1)\}$ where $(i,j) \in \{1,\ldots,k'\} \times \{1,\ldots,k'-1\}$, we define the graph

$$G_{i,j}^{\text{ver}} = G[\bigcup_{j' \in \{\alpha(j),\ldots,\alpha(j+1)\}} \phi^{-1}(\alpha(i),j')]$$

and we find the path $P_{i,j}^{\text{ver}}$ in it starting from $v_{i,j}$ finishing in $v_{i,j+1}$ and containing an edge $e_{i,j}^{\text{ver}}$ of $\sigma^{-1}(E(H))$ that belongs in $\phi^{-1}(\mathcal{E}(U_{i,j}^{\text{ver}})) \cup \phi^{-1}(U_{i,j}^{\text{ver}})$. As before, $P_{i,j}^{\text{ver}}$ is decomposed to a path $\downarrow P_{i,j}$ (containing $v_{i,j}$), the edge $e_{i,j}^{\text{ver}}$, and the path $\uparrow P_{i,j+1}$ (containing $v_{i,j+1}$). Let, finally, E^* be the set containing each $e_{x,y}^{\text{hor}}$ and each $e_{x,y}^{\text{ver}}$.

From Lemma 5, to prove that $L_{k'}$ is a minor of H, it is enough to define a function $\tau : E(G^\ell) \to V(L_{k'}) \cup E(L_{k'}) \cup \{\star\}$ certifying that $L_{k'}$ is a minor of G in a way that $\forall f \in E(L_{k'})$ $|\tau^{-1}(f) \cap \sigma^{-1}(E(H)| = 1$. For this, for every $(x,y) \in \{1,\ldots,k'\}$ we define $E_{x,y}$ as the union of the edges and the loops of the vertices of every path that exists in the set $\{\overleftarrow{P}_{x,y}, \overrightarrow{P}_{x,y}, \downarrow P_{x,y}, \uparrow P_{x,y}\}$ and for each $e \in E_{x,y}$ we set $\tau(e) = (x,y)$. Notice that for every $(x,y) \in \{1,\ldots,k'\}^2$, $G[\tau^{-1}(x,y)]$ is the union of a set of paths of G having a vertex in common, thus it induces a connected subgraph of G. Let now e be an edge of $L_{k'}$. In case $e = \{(x,y),(x+1,y)\}$ (resp. $e = \{(x,y),(x,y+1)\}$), then, by its definition, the edge $e_{x,y}^{\text{hor}}$ (reps. $e_{x,y}^{\text{ver}}$) connects an endpoint v_1 of an edge in $\tau^{-1}(x,y)$ (resp. $\tau^{-1}(x,y)$) with an endpoint v_2 of an edge in $\tau^{-1}(x+1,y)$ (resp. $\tau^{-1}(x,y+1)$). In any case, we set $\tau(v_1 v_2) = e$. It follows that $\tau(E^*) = E(L_{k'})$. For all edges of G^ℓ whose image has not been defined so far, we set $\tau(e) = \star$. It is now easy to verify that τ is a well-defined function and that $L_{k'}$ is a τ-generated minor of G.

Next we prove that $\forall f \in E(L_{k'})$ $|\tau^{-1}(f) \cap \sigma^{-1}(E(H)| = 1$. For this, first of all notice that, by the definition of τ, all edges in $\tau^{-1}(E(L_{k'})) = E^*$ are edges of $\sigma^{-1}(E(H))$. Therefore, it suffices to prove that for each $e \in E(H)$, $\sigma^{-1}(e)$ contains no more than one edge from E^*. Suppose in contrary that $e_1, e_2 \in \sigma^{-1}(e) \cap E^*$ and $e_1 \neq e_2$. As $\sigma(e_1) = \sigma(e_2) = e$, it follows that each e_i has an endpoint w_i in $\sigma^{-1}(w)$ and an endpoint z_i in $\sigma^{-1}(z)$, where $wz = e$. Since each subgraph $G[\sigma^{-1}(w)]$ and $G[\sigma^{-1}(z)]$ is connected, has at most c edges and both are disjoint, there are two disjoint paths of length at most c in G from w_1 to w_2 and from z_1 to z_2, a contradiction to $(*)$ as $e_1, e_2 \in E^*$. $\qquad \square$

Lemma 7. *Let H_1 and H_2 be two graphs. Consider a graph G such that H_1 is a c_1-contraction of G and H_2 is a c_2-contraction of G. If H_1 is planar then $\mathbf{tw}(H_2) = O(c_1 \cdot c_2 \cdot \mathbf{bg}(H_2)) = 36 \cdot (c_1 + 1) \cdot (c_2 + 1) \cdot [\mathbf{bg}(H_2) - 1] + O(c_1)$.*

Proof. Let H_1, H_2 be two contractions of G generated by some $\sigma_i : E(G^\ell) \to V(H_i) \cup E(H_i)$, $i = 1, 2$ respectively. Let $r = \mathbf{tw}(H_2)$. As G contains H_2 as a contraction, it follows that $\mathbf{tw}(G) \geq r$. By Lemma 3, $\mathbf{tw}(H_1) \geq (r+1)/(c_1+1) - 1$. Since H_1 is planar, by Lemma 2, H_1 contains $L_{r'}$ as a distance minor, where $r' = \lfloor \frac{1}{18} \cdot (\frac{r+1}{c_1+1} - 1) \rfloor$.

As H_1 is a contraction of G, then also G contains $L_{r'}$ as a distance minor. By Lemma 6, H_2 contains as a minor an (r'', r'')-grid, where $r'' = \lfloor \frac{r'-1}{2(c_2+1)} \rfloor + 1$, as claimed. $\qquad\square$

Proof of Theorem 1. Given a planar drawing of the lines of $\mathcal{B}$, consider any crossing p, that is a point of the plane that belongs to more than one line. By the assumptions, p belongs to exactly two lines, say L_1, L_2 and there is an open disc D of the plane containing p, but no lines other than L_1 and L_2 and no other point that belongs to more than one line. In addition, w.l.o.g. we can always assume that p is not an endpoint of L_1 or L_2; or else we can stretch inside D the line that ends in p without further altering the setting.

Then, let G_1 be the simple graph with an embedding in the plane, in which all endpoints of lines in $\mathcal{B}$ are vertices of G_1 and every line $L \in \mathcal{B}$ is an edge of G_1 joining the two vertices, which are endpoints of L. Note that the graph G_1 is not necessarily planar – in fact, any crossing of two lines in $\mathcal{B}$ is as well a crossing of the corresponding edges of G_1 in this embedding.

For every crossing p of two lines L_1, L_2 in $\mathcal{B}$ and hence of the corresponding edges e_1, e_2 of G_1, we can consider as above an open disc D of the plane in a way, so that $D \cap e = \emptyset$ for any edge $e \in E(G_1) \setminus \{e_1, e_2\}$, the only point in D that belongs to both edges is p, no vertex of G_1 lies in D and all considered discs are pairwise disjoint. Then, we subdivide e_1 and e_2 by adding two new vertices x, y in $D \setminus \{p\}$ and we join x and y with a new edge f that lies entirely in the disc D and meets L_1 and L_2 only at its endpoints. We denote as M the set of these new edges. Notice that we can contract the edge f inside the disc D so that the resulting vertex is the point p, leaving the embedding of the graph outside D untouched. By doing so for every edge in M, we obtain a planar embedding of a graph. Let H be this graph and let G be the graph before contracting the edges in M, i.e. $G/M = H$. Clearly, H is an 1-contraction of G. Moreover, if we contract all edges of G that are not in M, we obtain the intersection graph $G_\mathcal{B}$. Since every edge of G_1 was subdivided into at most $\xi + 1$ edges of G, the graph $G_\mathcal{B}$ is a $(\xi + 1)$-contraction of G and the result follows from Lemma 7. $\qquad\square$

4 Modeling body intersections by intersection of polysegments

Let $\mathcal{B} = \{B_1, \ldots, B_k\}$ be a collection of 2-dimensional geometric bodies in the plane. We assume that if two bodies do intersect each other, then every connected component of the intersection has a non-empty interior. Our goal is to associate each geometric body $B \in \mathcal{B}$ with a polysegment C such that the resulting set $\mathcal{B}'$ of polysegments conveys all necessary information regarding the disposition of the bodies in the plane and their intersections.

For every body B_i let us pick a point of the sphere p_i that lies in B_i and for every body B_j touching B_i $(i \neq j)$ in $\mathcal{B}$, a point p_{ij} that lies in $B_i \cap B_j$. We can assume without loss of generality, that these points are pairwise distinct and that any three of

them are not co-linear. We stress that since $\mathcal{B}$ is finite this assumption is safe, because we can always consider an open disc D_p of small radius around any given point p of the sphere, such that if p lies in a body B then we can replace B with the possibly expanded body $B' = B \cup D_p$ without altering the intersection graph of $\mathcal{B}$. Let now $\mathcal{P}_i$ be the set of all points that contain the index i in the assigned subscript.

A geometric body B of the sphere is ρ-convex, if for any two points of B there exists a polysegment of length ρ that lies entirely inside B and its endpoints are the given two points. Notice, that the definition of a ρ-convex body naturally extends the standard definition of a convex body, which under this new perspective is also called 1-convex.

Lemma 8. *For any collection of ρ-convex bodies $\mathcal{B}$ on the sphere, there exists a collection of polysegments $\mathcal{B}'$ and a bijection $\phi : \mathcal{B} \to \mathcal{B}'$, such that two bodies in $\mathcal{B}$ touch if and only if the corresponding polysegments in $\mathcal{B}'$ touch. Moreover, each polysegment $C \in \mathcal{B}'$ is crossed by the polysegments from $\mathcal{B}' \setminus C$ at most $\xi = O(\rho^2 \Delta^3)$ times, where Δ is the maximum degree in the intersection graph $G_\mathcal{B}$ of $\mathcal{B}$.*

Proof. Let us then consider such a collection $\mathcal{B}$ of ρ-convex bodies $\{B_1, \ldots, B_k\}$ where any body in $\mathcal{B}$ touches at most Δ other bodies, for some positive integers ρ and Δ. Clearly $|\mathcal{P}_i| \leq \Delta + 1$, for $1 \leq i \leq k$. Notice that as a body B_i of $\mathcal{B}$ is ρ-convex, for any two points in $\mathcal{P}_i$ there is a polysegment of length at most ρ that joins them. We create the following drawing of a polysegment.

Pick two points in $\mathcal{P}_i$ and join them with a polysegment that lies in B_i. As long as there are still points in $\mathcal{P}_i$ we did not join, pick one and join it to a $\mathcal{P}_i$ point or bend point of the so far drawing, so that the new polysegment does not cross the drawing. To see this is possible, consider that if the polysegment crosses a straight line segment of the drawing so far, we can simply replace it by one that joins the new point of $\mathcal{P}_i$ to a bend point that lies on the border of the straight line segment in question. In the end we obtain the drawing of a tree on the sphere containing all points in $\mathcal{P}_i$. At each step we joined at least one new point in $\mathcal{P}_i$, and in doing so we introduced at most $\rho - 1$ new bend points. It follows that the drawing has in total at most $\rho \cdot \Delta + 1$ bend and P_i points.

Finally, we circumscribe around the drawing of this tree a polysegment C_i which contains all points in $\mathcal{P}_i$, so that to every straight line segment of the tree correspond two straight line segments of C_i, and any bend point of C_i lies in an open disc of small radius around the corresponding bend point of the drawing of the tree. Clearly the polysegment C_i has length at most $2\rho \cdot \Delta$.

We have thus showed, that for any ρ-convex body $B_i \in \mathcal{B}$, there exists a polysegment C_i of length $O(\rho \cdot \Delta)$, that lies entirely inside B_i and contains all points in $\mathcal{P}_i$. The mapping of a body B_i to a polysegment C_i defines a bijection ϕ between the elements of the collection $\mathcal{B}$ and the collection of polysegments $\mathcal{B}' = \{C_i : 1 \leq i \leq k\}$. By construction, two distinct polysegments C_i, C_j both contain the point p_{ij} and hence they do share a point, if the corresponding bodies B_i, B_j touch. On the other hand,

as the polysegments lie inside the geometrical bodies, the polysegments C_i, C_j share a point of the sphere only if the associated bodies touch. This also implies that a polysegment C_i shares at most $O((\rho \cdot \Delta)^2 \cdot \Delta) = O(\rho^2 \cdot \Delta^3)$ points of the sphere with other polysegments in $\mathcal{B}'$, concluding the proof. $\qquad\square$

Straightforwardly applying Theorem 1 to the sets of polysegments constructed in Lemma 8, results to the following theorem for ρ-convex geometric bodies.

Theorem 2. *Let $\mathcal{B}$ be a set of ρ-convex bodies such that for each $B_1, B_2 \in \mathcal{B}$ with $B_1 \cap B_2 \neq \emptyset$, the set $B_1 \cap B_2$ has non-empty interior. Let $G_\mathcal{B}$ be the intersection graph of $\mathcal{B}$ and let Δ be the maximum degree of $G_\mathcal{B}$. Then $\mathbf{tw}(G_\mathcal{B}) = O(\rho^2 \Delta^3 \cdot \mathbf{bg}(G_\mathcal{B}))$.*

Given a positive real number α, we define the class of α-*fat convex intersection graphs* as the class containing an intersection graph $G_\mathcal{B}$ of a collection $\mathcal{B}$ of convex bodies, if the ratio between the maximum and the minimum radius of a circle where all objects in $\mathcal{B}$ can be circumscribed, and inscribed respectivelly, is upper bounded by α. The following lemma describes the manner in which the convex bodies of such a collection are being modeled by polysegments.

Lemma 9. *Let H be a graph on h vertices and let $\mathcal{B}$ be a collection of convex bodies on the sphere. If the intersection graph of $\mathcal{B}$ is α-fat and does not contain graph H as a subgraph, then there exists a collection of polysegments $\mathcal{B}'$ and a bijection $\phi : \mathcal{B} \to \mathcal{B}'$ such that two bodies in $\mathcal{B}$ touch if and only if the corresponding polysegments in $\mathcal{B}'$ touch. Moreover, each polysegment $C \in \mathcal{B}'$ is crossed by the polysegments from $\mathcal{B}' \setminus C$ at most $\xi = O(\alpha^6 \cdot h^3)$ times.*

Proof. Let us then consider such a collection $\mathcal{B}$ of convex bodies $\{B_1, \ldots, B_k\}$ where the intersection graph $G_\mathcal{B}$ is α-fat and has maximum degree Δ, for a positive integer Δ and a positive real α. It follows that for a set of points $\mathcal{P}_i$ corresponding to a body B_i of $\mathcal{B}$, it holds that $|\mathcal{P}_i| \leq \Delta + 1$. As the body B_i is convex, there exists a polysegment C_i of length $|\mathcal{P}_i| - 1$ containing all points in $\mathcal{P}_i$, that lies entirely in B_i. Let $\mathcal{B}'$ be the collection of the polysegments $\{C_i : 1 \leq i \leq k\}$. Clearly, two distinct polysegments of $\mathcal{B}'$ share a point of the sphere if and only if the corresponding bodies touch. Furthermore, a polysegment has at most $O(\Delta^3)$ common points with other polsegments in $\mathcal{B}'$.

We claim that $\Delta \leq 16\alpha^2 \cdot h$, which directly implies the bound of the lemma. To contradiction, assume that $G_\mathcal{B}$ has a vertex with degree more than $16h\alpha^2$. Then, for the corresponding body B of the collection $\mathcal{B}$, pick an arbitrary point $x \in B$. Since each object in $\mathcal{B}$ can be inscribed in a circle of radius R, all bodies intersecting B belong to the ball C of radius $4R$ with the center in x. This ball has area $16\pi R^2$. Since each object in $\mathcal{B}$ contains a circle of radius r, the sum of the areas of the bodies intersecting B exceeds $16h\alpha^2 \times \pi r^2 = 16\pi h R^2$. Therefore, there is a point in the ball C belonging to at least $h + 1$ bodies from $\mathcal{B}$. Thus, $G_\mathcal{B}$ contains a clique of size $h + 1$ that contradicts to the assumption that it is H-free. $\qquad\square$

Again, by straightforwardly applying Theorem 1 to the sets of polysegments constructed in Lemma 9, we derive an improved theorem for H-free α-fat convex intersection graphs of geometric bodies.

Theorem 3. *Let H be a graph on h vertices, and let $\mathcal{B}$ be a collection of convex bodies on the sphere such that for each $B_1, B_2 \in \mathcal{B}$ with $B_1 \cap B_2 \neq \emptyset$, the set $B_1 \cap B_2$ has a non-empty interior. If the intersection graph $G_{\mathcal{B}}$ of $\mathcal{B}$ is α-fat and does not contain H as a subgraph, then $\mathbf{tw}(G_{\mathcal{B}}) = O(\alpha^6 h^3 \cdot \mathbf{bg}(G_{\mathcal{B}}))$.*

5 Conclusions and further research

We believe that the applicability of our combinatorial results is even wider than what is explained in the previous section. The main combinatorial engine of this paper is Lemma 7 that essentially induces an edit-distance notion between graphs under contractibility. This is materialized by the following definition.

Definition 2. *Let G_1 and G_2 be graphs. We define the* contraction-edit distance *between G_1 and G_2, denoted by $\mathbf{cdist}(G_1, G_2)$, as the minimum c for which there exists a graph that contains both G_1 and G_2 as c-contractions. Given a graph G we define $\mathcal{B}^c(G) = \{H \mid \mathbf{cdist}(G, H) \leq c\}$. Finally, given a graph class $\mathcal{G}$, we define $\mathcal{B}^c(\mathcal{G}) = \bigcup_{G \in \mathcal{G}} \mathcal{B}^c(G)$. We refer to the class $\mathcal{B}^c(\mathcal{G})$ as the c-contraction extension of the class $\mathcal{G}$.*

A direct consequence of Lemma 7 is the following:

Corollary 1. *Let $\mathcal{P}$ be the class of planar graphs, then for every fixed constant c, $\mathcal{B}^c(\mathcal{P})$ satisfies* (1).

Actually, Corollary 1 can be extended much further than planar graphs. For this, the only we need analogues of Lemma 2 for more general graph classes. Using the main result of [7], it follows that Lemma 2 is qualitatively correct for every graph class that excludes an apex graph as a minor (an apex graph is a graph that can become planar after the removal of a vertex). By plugging this more general version of Lemma 2 to the proofs of the previous section we obtain the following.

Theorem 4. *Let H be an apex-minor free graph and let $\mathcal{G}_H$ be the class of graphs excluding H as a minor. Then, for every fixed constant c, the class $\mathcal{B}^c(\mathcal{G}_H)$ satisfies* (1).

All the algorithmic applications of this paper follow by the fact that all geometric intersection graph classes considered in this paper are subsets of $\mathcal{B}^c(\mathcal{P})$ for some choice of c. Clearly, Theorem 4 offers a much more wide framework for this, including graphs of bounded genus (including intersection graphs of lines or polygons on surfaces), graphs excluding a single-crossing graph, and $K_{3,r}$-minor free graphs. We believe that Theorem 4, that is the most general combinatorial extension of our results may have applications to more general combinatorial objects than just intersection graph classes. We leave this question open for further research.

Acknowledgements

References

[1] H. L. Bodlaender, R. G. Downey, M. R. Fellows, and D. Hermelin. On problems without polynomial kernels. *J. Comput. Syst. Sci.*, 75:423–434, December 2009.

[2] L. Cai and D. Juedes. On the existence of subexponential parameterized algorithms. *J. Comput. System Sci.*, 67(4):789 – 807, 2003.

[3] E. Demaine and M. Hajiaghayi. The bidimensionality theory and its algorithmic applications. *The Computer Journal*, 51(3):292–302, 2007.

[4] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Subexponential parameterized algorithms on graphs of bounded genus and H-minor-free graphs. *Journal of the ACM*, 52(6):866–893, 2005.

[5] E. D. Demaine and M. Hajiaghayi. Bidimensionality: new connections between FPT algorithms and PTASs. In *Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 590–601 (electronic). ACM, New York, 2005.

[6] E. D. Demaine and M. Hajiaghayi. Linearity of grid minors in treewidth with applications through bidimensionality. *Combinatorica*, 28(1):19–36, 2008.

[7] F. V. Fomin, P. A. Golovach, and D. M. Thilikos. Contraction obstructions for treewidth. *J. Comb. Theory, Ser. B*, 101(5):302–314, 2011.

[8] F. V. Fomin, D. Lokshtanov, V. Raman, and S. Saurabh. Bidimensionality and EPTAS. In *22st ACM–SIAM Symposium on Discrete Algorithms (SODA 2011)*, pages 748–759. ACM-SIAM, San Francisco, California, 2011.

[9] F. V. Fomin, D. Lokshtanov, and S. Saurabh. Bidimensionality and geometric graphs. In *23st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)*, pages 1563–1575, 2012.

[10] F. V. Fomin, D. Lokshtanov, S. Saurabh, and D. M. Thilikos. Bidimensionality and kernels. In *21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2010)*, pages 503–510, 2010.

[11] Q.-P. Gu and H. Tamaki. Improved bounds on the planar branchwidth with respect to the largest grid minor size. *Algorithmica*, 64(3):416–453, 2012.