Geometric Extensions of Cutwidth in any Dimension*

Menelaos I. Karavelas† Spyridon Maniatis‡ Dimitrios M. Thilikos‡
Dimitris Zoros‡

Abstract

We define a multi-dimensional geometric extension of cutwidth. A graph has d-cutwidth at most k if it can be embedded in the d-dimensional euclidean space so that no hyperplane can intersect more than k of its edges. We prove a series of combinatorial results on d-cutwidth which imply that for every d and k, there is a linear time algorithm checking whether the d-cutwidth of a graph G is at most k.

1 Introduction

The cutwidth of a (total) vertex ordering of a graph is the maximum number of edges connecting vertices on opposite sides of any of the "gaps" between successive vertices in the linear layout. The cutwidth of a graph G, denoted by $\text{cw}(G)$, is the minimum cutwidth over all its possible vertex orderings. The problem that asks, given a n-vertex graph G and an integer k, whether $\text{cw}(G) \leq k$, is an NP-complete problem known in the literature as the Minimum Cut Linear Arrangement problem [4]. From the parameterized complexity point of view, the same problem is fixed parameter tractable, as an algorithm that checks whether cutwidth$(G) \leq k$ in $f(k) \cdot n$ steps was given in [10]. Cutwidth has been extensively studied both from its combinatorial (see e.g. [2, 7, 1]) as well as its algorithmic point of view [8, 11, 3, 6].

d-dimensional cutwidth. In this note we introduce a multi-dimensional geometric extension of cutwidth, namely the d-dimensional cutwidth (or, simply, d-cutwidth) that, roughly, instead of mono-dimensional linear arrangements of the graph G, we consider embeddings of G in the d-dimensional Euclidean space $\mathbb{R}^d$ and define the d-cutwidth of such an embedding to be the maximum number of edges that can be intersected by a hyperplane of $\mathbb{R}^d$. Then, the d-cutwidth of G, denoted by $\text{cw}_d(G)$, is the minimum d-cutwidth over all such embeddings. Our results are summarized in the following.

Theorem 1 *The following hold:*

*The last three authors were co-financed by the E.U. (European Social Fund - ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) - Research Funding Program: "Thales. Investing in knowledge society through the European Social Fund". Emails: `mkaravel@iacm.forth.gr`, `spyridon.maniatis@gmail.com`, `sedthilk@thilikos.info`, `dzoros@math.uoa.gr`.

†Dept. of Mathematis and Applied Mathematics, University of Crete, Heraklion, Greece.

‡Dept. of Mathematics, National & Kapodistrian University of Athens, Athens, Greece.

i. d-cutwidth is immersion closed[1].
ii. For every graph G and every $d \geq 1$, $\mathrm{cw}_d(G) \leq \mathrm{cw}_{d+1}(G)$.
iii. For every graph G and every $d \geq 1$, $\mathrm{cw}_d(G) \leq d \cdot \mathrm{cw}(G)$.
iv. For every graph G, $\mathrm{cw}_3(G) \leq 2 \cdot \mathrm{cw}_2(G)$.

2 Preliminaries and definitions

Hyperplanes and hyperspheres. Every $(d-1)$-dimensional subspace Π of a d-dimensional space $\mathcal{X}$ is called a *hyperplane* of $\mathcal{X}$. Here we are interested in hyperplanes of $\mathbb{R}^d$ (subspaces isomorphic to $\mathbb{R}^{d-1}$). Let Π be a hyperplane in $\mathbb{R}^d$, then there are $a_0, a_1, \ldots, a_d \in \mathbb{R}$ such that $\Pi = \{(x_1, \ldots, x_d) \in \mathbb{R}^d \mid a_1x_1 + \cdots + a_dx_d + a_0 = 0\}$. We denote by $H(d)$ the set of all hyperplanes of $\mathbb{R}^d$. A *hypersphere*, $S(c, r)$, with *center* c and *radius* r in $\mathbb{R}^d$ is the set $\{(x_1, \ldots, x_d) \in \mathbb{R}^d \mid \sum_{i=1}^{d}(x_i - c_i)^2 = r^2\}$. We denote by $S(d)$ the set of all hyperspheres of $\mathbb{R}^d$.

Curves of $\mathbb{R}^d$. We call a continuous function $C : [0, 1] \to \mathbb{R}^d$ a *curve* of $\mathbb{R}^d$ with ends $C(0)$ and $C(1)$.

Graph Embeddings in $\mathbb{R}^d$. Let $G = (V, E)$ be a graph. An embedding of G, denoted by $\mathcal{E}_d(G)$, in the euclidean space $\mathbb{R}^d$ is a tuple $(f, \mathcal{C})$, where $f : V \to \mathbb{R}^d$ is an injection, mapping the vertices of G to $\mathbb{R}^d$ and $\mathcal{C} = \{C_e \mid e \in E\}$ is a set of curves of $\mathbb{R}^d$ with the following properties:

1. for every $e = \{u, v\} \in E$, the ends of C_e are $f(u)$ and $f(v)$
2. for all $x \in (0, 1)$ and for all $v \in V$ it holds that $C_e(x) \neq f(v)$.

For simplicity, we may sometimes refer to the elements of $f(V)$ and $\mathcal{C}$ as the vertices and edges of $\mathcal{E}_d(G)$ respectively.

Essential embeddings, $\mathbf{E}_d(G)$. We denote by $\mathbf{E}_d(G)$ the set of all embeddings $\mathcal{E}_d(G) = (f, \mathcal{C})$, of G in $\mathbb{R}^d$, such that for every positive integer $i \leq d$, if S is a subset of V with $|S| \geq i$, then the dimension of the subspace defined by $\{f(u) \mid u \in S\}$ is $i - 1$. We call every element of $\mathbf{E}_d(G)$ an *essential-embedding of G in $\mathbb{R}^d$*. Let $\mathcal{E}_d(G)$ be an essential-embedding of G in $\mathbb{R}^d$, then if Π is a hyperplane of $\mathbb{R}^d$ (resp. Σ is a hypersphere of $\mathbb{R}^d$) that does not intersect any $f(v)$, $v \in V$, we denote by $\partial_G(\mathcal{E}_d(G), \Pi)$ (resp. $\partial_G(\mathcal{E}_d(G), \Sigma)$) the set of curves of $\mathcal{E}_d(G)$ that are intersected by Π (resp. Σ).

Vertical projection. Given a point $x = (x_1, \ldots, x_d)$ in $\mathbb{R}^d$ and a hyperplane $\Pi = \{(x_1, \ldots, x_d) \in \mathbb{R}^d \mid a_1x_1 + \cdots + a_dx_d + a_0 = 0\}$, the vertical projection of x on Π is the point $y = \{y_1, \ldots, y_d\} \in \mathbb{R}^d$ where

$$y_i = x_i - a_i \frac{a_1x_1 + \cdots + a_dx_d + a_0}{a_1{}^2 + \cdots + a_d{}^2},\ i \in \{1, \ldots, d\}.$$

[1]A graph H is an *immersion* of a graph G if it can be obtained from G after a sequence of vertex/edge removals or edge lifts (the operation of *lifting* two edges $\{x, y\}$ and $\{y, z\}$ incident to the same vertex y is the operation of replacing these edges by the edge $\{x, z\}$). A graph invariant is *immersion closed* if its value on a graph G is always smaller or equal than its value on its immersions.

Definition 1 *Let $G = (V, E)$ be a graph and k, d be positive integers, where $d \geq 2$. Then, we define the d-dimensional cutwidth of G, or simply d-cutwidth, to be*

$$\operatorname{cw}_d(G) = \min_{\mathcal{E}_d(G) \in \mathbf{E}_d(G)} \max\{|\partial_G(\mathcal{E}_d(G), \Pi)| \mid \Pi \in H(d)\}$$

We say that an embedding $\mathcal{E}_d(G) \in \mathbf{E}_d(G)$ *realizes* d-cutwidth of G if for every hyperplane Π of $\mathbb{R}_d$, $|\partial_G(\mathcal{E}_d(G), \Pi)| \leq \operatorname{cw}_d(G)$ and the equation holds for at least one hyperplane. Observe that any hyperplane Π of $\mathbb{R}_d$ that meets a curve $C_e \in \mathcal{C}$ once (if it meets a curve more than once it is easy to observe that this particular embedding does not realize the d-cutwidth of G), also meets the unique straight line segment of $\mathbb{R}^d$ with parametric equation $\sigma_e(t) = t \cdot C_e(0) + (1-t) \cdot C_e(1)$, $t \in \mathbb{R}$, i.e., the straight line segment of $\mathbb{R}^d$ that is defined by the "images" of the endpoints of edge e. Therefore, without loss of generality, we can consider only *straight-line embeddings* where $\mathcal{C} = \{\sigma_e \mid e \in E\}$. Notice that every straight line embeding $\mathcal{E}_d(G) = (f, \mathcal{C})$ is fully defined by the function f, therefore, for simplicity, from now on we will omit $\mathcal{C}$.

Cutwidth. The *cutwidth* of a graph $G = (V, E)$ with n vertices is defined as follows. Let $L = \langle v_1, \ldots, v_n \rangle$ be a layout of V. For $i = 1, \ldots, n-1$, we define the *cut at position* i, denoted by $\partial G(L, i)$, as the set of crossover edges of G that have one endpoint in $\{v_1, \ldots, v_i\}$ and one in $\{v_{i+1}, \ldots, v_n\}$. The *width of a layout* L of $V(G)$ is equal to $\max\{|\partial_G(L, i)| \mid 1 \leq i \leq n-1\}$ and the *cutwidth* of G is the minimum width over all the orderings of $V(G)$.

According to the notation of Definition 1 we can give an equivalent definition of cutwidth as follows:

Let $G = (V, E)$ be a graph. An embedding of G in $\mathbb{R}$, denoted $\mathcal{E}_1(G)$, is a tuple $(f, \mathcal{I})$, where $f : V \to \mathbb{R}$ is an injection, mapping the vertices of G to $\mathbb{R}$ and $\mathcal{I} = \{(f(u), f(v)) \subset \mathbb{R} \mid \{u, v\} \in E\}$ is a set of open intervals of $\mathbb{R}$. Given an embedding $\mathcal{E}_1(G) = (f, \mathcal{I})$, we denote by $\partial_G(\mathcal{E}_1(G), x)$ the set of intervals of $\mathcal{I}$ in which x belongs.

Definition 2 *Let $G = (V, E)$ be a graph and k a positive integer. We define 1-cutwidth, or simply cutwidth, of G to be*

$$\operatorname{cw}_1(G) = \min_{\mathcal{E}_1(G)} \max\{|\partial_G(\mathcal{E}_1(G), x)| \mid x \in \mathbb{R}\}.$$

Observe that the above definition of cutwidth is equivalent to its usual definition. Also observe that in this case, hyperplanes degenerate to subspaces of $\mathbb{R}$ of dimension 1, i.e. points, and our demand of essential embeddings is expressed by our demand of injective functions. Therefore, d-cutwidth is the intuitive generalization of the notion of cutwidth in any dimension $d \geq 2$.

3 Properties of d-cutwidth

In this section we will prove some properties of d-dimensional cutwidth.

Lemma 1 *For every graph G and every $d \geq 1$,*

$$\mathrm{CW}_d(G) \leq \mathrm{CW}_{d+1}(G).$$

Proof. Let $G = (V, E)$ be a graph, and let $\mathcal{E}_{d+1}(G) = f$ be an embedding of G in $\mathbb{R}^{d+1}$ that realizes $\mathrm{CW}_{d+1}(G)$. Let Π_0 be a hyperplane of $\mathbb{R}^{d+1}$ such that, for every $e \in E$, Π_0 is not vertical to σ_e. We vertically project $\mathcal{E}_{d+1}(G)$ on Π_0, which gives us an embedding $\mathcal{E}_d(G)$ of G in $\mathbb{R}^d$, as the restriction of f in $\mathbb{R}^d$ satisfies the conditions of a graph embedding.

Assume that there exists a hyperplane Π in $\mathbb{R}^d$ that intersects $\mathcal{E}_d(G)$ more than $\mathrm{CW}_{d+1}(G)$ times. Then we can construct a new hyperplane Π' in $\mathbb{R}^{d+1}$ that intersects $\mathcal{E}_{d+1}(G)$ more than $\mathrm{CW}_{d+1}(G)$ times. This hyperplane Π' is the hyperplane vertical to Π that passes through Π. But this fact leads to a contradiction, since any hyperplane in $\mathbb{R}^{d+1}$ can intersect $\mathcal{E}_{d+1}(G)$ at most $\mathrm{CW}_{d+1}(G)$ times. Hence our assumption that Π' exists is false, i.e., every hyperplane in $\mathbb{R}^d$ intersects $\mathcal{E}_d(G)$ at most $\mathrm{CW}_{d+1}(G)$ times. Therefore, $\mathrm{CW}_d(G) \leq \mathrm{CW}_{d+1}(G)$. $\square$

Proposition 1 *For every $d \geq 1$, $\mathrm{CW}_d(P_d) = d$, where by P_n we denote the path of length n.*

Proof. Let $\mathcal{E}_d(P_d)$ be an essential straight line-embedding of $P_d = (V, E)$ in $\mathbb{R}^d$, where $V = \{v_1, \ldots, v_{d+1}\}$ and $E = \{\{v_i, v_{i+1}\} \mid i \in \{1, \ldots, d-1\}\}$. Consider the midpoints m_i of $\sigma_{\{i,i+1\}}$, for every $i \in \{1, \ldots, d-1\}$. These d points of $\mathbb{R}^d$ define a hyperplane of $\mathbb{R}^d$ that intersects all edges of $\mathcal{E}_d(P_d)$. Thus $\mathrm{CW}_d(P_d) \geq d$ and as $|E(P_d)| = d$ we derive that $\mathrm{CW}_d(P_d) = d$. $\square$

Lemma 2 *For every graph G and every $d \geq 2$ we have:*

$$\mathrm{CW}_d(G) \leq d \cdot \mathrm{CW}(G).$$

Proof. Consider the d-dimensional curve C with parametric equation

$$C(t) := (t, t^2, t^3, \ldots, t^d), \quad t \in \mathbb{R}.$$

Consider an ordering of the nodes of G that realizes the cutwidth of G. Embed a node v_i of G to the point $p_i = C(t_i)$, for an appropriate value t_i. By appropriate we mean that if a node v_i is after a node v_j in the cutwidth ordering, then the parametric value t_i corresponding to v_i is strictly greater than the parameter value t_j corresponding to node v_j. Now embed an edge $e_{ij} = (v_i, v_j)$ of G by connecting the points p_i and p_j on C with the minimum length arc of C connecting these points.

Consider a generic hyperplane Π with equation $a_1x_1 + a_2x_2 + \ldots + a_dx_d + a_0 = 0$, where, for all i, $a_i \in \mathbb{R}$. Π can cut C at at most d points. To see that, solve the system of equations

$$\begin{aligned} a_1x_1 + a_2x_2 + \ldots + a_dx_d + a_0 &= 0, \\ x_i &= t^i, \quad i = 1, \ldots, d \end{aligned}$$

for t. This gives the polynomial equation $q(t) := a_0 + a_1 t + a_2 t^2 + \ldots + a_d t^d = 0$, in t of maximum degree d. Since $q(t) = 0$ has at most d real roots, we deduce that Π intersects C at at most d points. At each point of intersection at most $\text{CW}(G)$ edges of the embedding of G pass through that point. Hence, Π intersects at most $d \cdot \text{CW}(G)$ edges of G, i.e., $\text{CW}_d(G) \leq d\ \text{CW}(G)$. □

Corollary 1 *For $G = P_d$ the inequality of Lemma 2 becomes an equation. Hence, Lemma's 2 bound is tight.*

Proof. We have from Proposition 1, that for every $d \geq 1$, $\text{CW}_d(P_d) = d$, thus $\text{CW}_d(P_d) = d \cdot \text{CW}_1(P_d)$. This proves that the inequality of Lemma 2 can be tight.

Definition 3 *Let $G = (V, E)$ be a graph and k, d be positive integers, where $d \geq 2$. Then we define the spherical d-dimensional cutwidth of G, or simply spherical d-cutwidth, to be*

$$\text{SCW}_d(G) = \min_{\mathcal{E}_d(G) \in \mathbf{E}_d(G)} \max\{|\partial_G(\mathcal{E}_d(G), \Sigma)| \mid \Sigma \in S(d)\}$$

Lemma 3 *For every graph G and any $d \geq 2$ we have:*

$$\text{CW}_d(G) \leq \text{SCW}_d(G) \leq (d+1)\ \text{CW}(G).$$

Proof. The left-most inequality is obvious. For every embedding $\mathcal{E}_d(G)$ of G in $\mathbb{R}^d$, the number of intersections of $\mathcal{E}_d(G)$ with a generalized hypersphere (of an appropriate center and radius) in $\mathbb{R}^d$ is greater or equal to the number of intersections of $\mathcal{E}_d(G)$ with a hyperplane in $\mathbb{R}^d$. Hence, $\text{CW}_d(G) \leq \text{SCW}_d(G)$.

Now consider the curve $C(t) = (t, t^2, \ldots, t^d)$, with $t \in \mathbb{R}^d$, and consider the ordering of the nodes of G that realizes $\text{CW}(G)$. We embed the i-th node v_i of G, in this ordering, to the point $p_i = C(i)$. We embed an edge $e_{ij} = (v_i, v_j)$ of G by connecting the points p_i and p_j on C with the minimum length arc of C connecting these points.
We claim that this curve has at most $d+1$ intersections with a generalized hypersphere S in $\mathbb{R}^d$. If S is actually a plane then we can simply apply the argumentation presented in the proof of Lemma 2. Suppose now that S is a true hypersphere, and let $\sum_{i=1}^{d}(x_i - a_i)^2 + a_0 = 0$ be the equation of S, where, for all i, $a_i \in \mathbb{R}$. Consider the following system of equations:

$$\begin{aligned} \sum_{i=1}^{d}(x_i - a_i)^2 + a_0 &= 0, \\ x_i &= t^i, \quad i = 1, \ldots, d. \end{aligned}$$

The intersections of S with the afore-mentioned embedding of G in C is bounded by the number of real solutions of the system above, for which t is positive. Solving this system for t we get a polynomial equation for t, namely:

$$\sum_{i=1}^{d}(t^i - a_i)^2 + a_0 = 0.$$

Expanding the above equation we get:

$$\sum_{i=1}^{d} a^{2i} - 2\sum_{i=1}^{d} a_i t^i + \sum_{i=1}^{d} a_i^2 + a_0 = 0,$$

which can be rewritten as:

$$\sum_{i=1}^{d} t^{2i} - 2\sum_{i=1}^{\lfloor \frac{d}{2} \rfloor} a_{2i} t^{2i} - 2\sum_{i=1}^{\lceil \frac{d}{2} \rceil} a_{2i-1} t^{2i-1} + \sum_{i=1}^{d} a_i^2 + a_0 = 0.$$

It is fairly easy to verify that the above equation can actually be rewritten as:

$$\sum_{i=\lfloor \frac{d}{2} \rfloor+1}^{d} t^{2i} + \sum_{i=1}^{d} \left(\frac{1+(-1)^i}{2} - 2a_i \right) t^i + \sum_{i=1}^{d} a_i^2 + a_0 = 0.$$

By Descartes' rule of signs, the number of positive real roots of this polynomial is bounded above by the number of sign variations in the sequence of its (non-zero) coefficients. Taking a close look at this polynomial, we observe that its first $d - \lfloor \frac{d}{2} \rfloor$ non-zero coefficients are equal to 1, which implies that the number of sign variations in the sequence of its coefficients is fully determined by the last $d+2$ coefficients. A sequence of $d+2$ real numbers can have at most $d+1$ sign variations, hence the number of positive real roots of this polynomial is at most $d+1$.

To finalize the proof, since any hypersphere in $\mathbb{R}^d$ intersects with C at most $d+1$ times, we conclude that the maximum number of intersections of the embedding of G in C is at most $(d+1)\ \textsc{cw}(G)$. Therefore, $\textsc{scw}_d(G) \leq (d+1)\ \textsc{cw}(G)$. $\square$

Lemma 4 *For every graph G and every $d \geq 1$ we have:*

$$\textsc{cw}_{d+1}(G) \leq \textsc{scw}_d(G).$$

Proof. Consider a graph G and an embedding $\mathcal{E}_d(G)$ in $\mathbb{R}^d$ for which $\textsc{scw}_d(G)$ is attained. Let us identify $\mathbb{R}^d$ with the hyperplane $x_{d+1} = 0$ in $\mathbb{R}^{d+1}$ and consider the unit hypersphere $\mathbb{S}^d$ in $\mathbb{R}^{d+1}$ centered at the origin. Let $\Sigma : \mathbb{R}^d \to \mathbb{S}^d$ be the stereographic projection from $\mathbb{R}^d$ to the unit hypersphere $\mathbb{S}^d$ in $\mathbb{R}^{d+1}$, and define $\mathcal{E}_{d+1}(G)$ to be the image of $\mathcal{E}_d(G)$ through the stereographic projection Σ, i.e., $\mathcal{E}_{d+1}(G) = \Sigma(\mathcal{E}_d(G))$.

Assume that there exists a hyperplane Π in $\mathbb{R}^{d+1}$ that cuts $\mathcal{E}_{d+1}(G)$ more than $\textsc{scw}_d(G)$ times. Let S be the intersection of Π with $\mathbb{S}^d$ and let S' be the inverse image of S with respect to the stereographic projection, i.e., $S' = \Sigma^{-1}(S)$. Since S is a hypersphere lying on $\mathbb{S}^d$, S' is either a hyperplane or hypersphere in $\mathbb{R}^d$. Since the stereographic projection preserves intersections, we deduce that S' intersects $\mathcal{E}_d(G)$ more than $\textsc{scw}_d(G)$ times. But this contradicts the definition of $\mathcal{E}_d(G)$, which implies that our assumption that Π cuts $\mathcal{E}_{d+1}(G)$ more than $\textsc{scw}_d(G)$ times is false. Hence, we found an embedding of G in $\mathbb{R}^{d+1}$ for which the maximum number of intersections with any hyperplane in $\mathbb{R}^{d+1}$ is at most $\textsc{scw}_d(G)$. Therefore, $\textsc{cw}_{d+1}(G) \leq \textsc{scw}_d(G)$. $\square$

Lemma 5 *For any graph G,*

$$\text{CW}_3(G) \leq 2\ \text{CW}_2(G).$$

Proof. Let G be a graph and consider an ordering of the nodes of G that realizes $\text{CW}(G)$. Consider an axis-aligned ellipse E centered at the origin with its x-axis being greater than its y-axis (e.g., the ellipse $4x^2 + y^2 - 4 = 0$). Embed the nodes of G on the positive half of E, denoted as $E_{1/2}$, that is on the half-ellipse that lies on the positive halfplane with respect to the x-axis, in such a way so that their x-coordinates preserve the ordering. In other words, given two nodes v_i and v_j of G, such that v_i precedes v_j in the node ordering, then $x_i < x_j$. Let us call $\mathcal{E}_2(G)$ the above-mentioned embedding. Given a generalized circle C on the plane, C can cut the half-ellipse $E_{1/2}$ at most two times. Hence we found an embedding of G in $\mathbb{R}^2$ such that any generalized circle C in $\mathbb{R}^2$ intersects this embedding at most $2\ \text{CW}(G)$ times. In other words, $\text{SCW}_2(G) \leq 2\ \text{CW}(G)$. Using the results from Lemmas 1 and 4 we conclude:

$$\text{CW}_3(G) \leq \text{SCW}_2(G) \leq 2\ \text{CW}(G) \leq 2\ \text{CW}_2(G),$$

which is what we wanted to prove. □

4 Immersions

Immersions. We say that a graph H is an *immersion (strong immersion)* of a graph $G = (V, E)$, $H \preceq G$, if we can obtain H from a subgraph (induced subgraph resp.) of G by *lifting edges*, where an edges lift is the following operation: given two adjacent edges, say $\{v, u\}$ and $\{u, w\}$, of G, we delete these edges from E and add the edge $\{v, w\}$ (Iin the case where $\{v, w\}$ already was present this operation creates a multiple edge). We can also define the immersion relation as follows: H is an immersion of G if there is an injective mapping $f : V(H) \to V(G)$, such that for every edge $\{u, v\}$ of H, there is a path from $f(u)$ to $f(v)$ in G and for any two edges of H the corresponding paths in G are *edge-disjoint*, that is, they do not share common edges. Additionally, if these paths are internally disjoint from $f(V(H))$, then we say that H is a strong immersion of G.

Theorem 2 *Let $G = (V, E)$ be a graph and H be an immersion (strong immersion) of G. For every $d \geq 1$, $\text{CW}_d(H) \leq \text{CW}_d(G)$, i.e. d-cutwidth is an immersion closed parameter.*

Proof. Let $\mathcal{E}_d(G)$ be an an embedding of G in $\mathbb{R}^d$ that realizes $\text{CW}_d(G)$. Let H' be a subgraph of G. Then, given $\mathcal{E}_d(G) = (f, \mathcal{C})$, we define an embedding $\mathcal{E}_d(H') = (f_{H'}, \mathcal{C}_{H'})$ of H' in $\mathbb{R}^d$, where $f_{H'}$ is the restriction of f to $V(H')$ and $\mathcal{C}_{H'} = \{c_e \in \mathcal{C} \mid e \in E(H')\}$. Observe that every hyperplane of $\mathbb{R}^d$ that intersected $l \leq \text{CW}_d(G)$ edges of $\mathcal{E}_d(G)$ intersects at most l edges of $\mathcal{E}_d(H')$, therefore $\text{CW}_d(H) \leq \text{CW}_d(G)$.

We will next prove that if H is the result of one lift of edges $e_1 = \{u, v\}, e_2 = \{v, w\}$ of E, then $\text{CW}_d(H) \leq \text{CW}_d(G)$. Clearly, for any hyperplane R of $\mathbb{R}^d$ its corresponding numbers $|\partial_H(\mathcal{E}_d(H), R)|$ and $|\partial_G(\mathcal{E}_d(G), R)|$ can differentiate only due to the intersections of R with edges e, e_1 and e_2. Let $E(H) \ni e = \{u, w\}$ be the resulting edge of the lift. Let $\Pi \in H(d)$ be a hyperplane that does not intersect $f(v)$, $\forall v \in V$. If Π intersects e (more accurately, the straight line segment that represents e in the embedding in $\mathcal{E}_d(G)$), then Π separates $\mathbb{R}^d$ into two halfspaces (i.e. subspaces of dimension d), namely A and B. Assume, without loss generality, that $u \in A$ and $w \in B$. Then, as Π does not intersect $f(v)$, either $v \in A$ or $v \in B$, which means that either Π intersects e_1 or Π intersects e_2 in $\mathcal{E}_d(G)$. Therefore, $|\partial_H(\mathcal{E}_d(H), \Pi)| \leq |\partial_G(\mathcal{E}_d(G), \Pi)|$. The same holds trivially for the case that Π does not intersect e. Therefore, $\mathrm{cw}_d(H) \leq \mathrm{cw}_d(G)$. Summing up the above we get that, for every graph H that is the result of some (maybe none) vertex deletions, edge deletions and edge lifts of G, $\mathrm{cw}_d(H) \leq \mathrm{cw}_d(G)$, which is what we stated. Notice that the above proof implies the same relation between a graph G and a strong immersion H of G. □

5 Algorithmic remarks about d-cutwidth

As a consequence of the result in [9], for every k, the class of immersion minimal graphs with d-cutwidth bigger than k contains a finite set of graphs. We call this class *immersion obstruction set* for cutwidth at most k and we denote it by $\mathcal{O}_k$. This fact, combined with Theorem 1.i, implies that $\mathrm{cw}_d(G) \leq k$ if and only if none of the graphs in $\mathcal{O}_k$ is contained in G as an immersion. According to the result of Grohe, Kawarabayashi, Marx, and Wollan in [5], checking whether an n-vertex graph contains as an immersion some k-vertex graph H, can be done in $f(k) \cdot n^3$ steps. As a consequence, checking whether $\mathrm{cw}_d(G) \leq k$ can be done in $f(k) \cdot n^3$ steps. This running time can become linear (on n) using the first inequality of Theorem 1.iii. Indeed, the algorithm first checks whether $\mathrm{cw}(G) \leq k$. If the answer is negative then we can safely report that $\mathrm{cw}_d(G) > k$. If not, then it is known (see e.g. [10]) that G has a tree decomposition of width $\leq k$ and to check whether some of the graphs in $\mathcal{O}_k$ is contained in G as an immersion can be done using dynamic programming in $f(k) \cdot n$ steps.

Unfortunately, the above algorithm is non-constructive as we have no other knowledge about the set $\mathcal{O}_k$, except from the fact that it is finite. To obtain a constructive $f(k) \cdot n$ step algorithm for d-cutwidth remains an insisting open problem.

References

[1] F. R. K. Chung and Paul D. Seymour. Graphs with small bandwidth and cutwidth. *Discrete Mathematics*, 75(1-3):113–119, 1989.

[2] Fan R. K. Chung. On the cutwidth and the topological bandwidth of a tree. *SIAM Journal on Algebraic and Discrete Methods*, 6(2):268–277, 1985.

[3] Moon Jung Chung, Fillia Makedon, Ivan Hal Sudborough, and Jonathan Turner. Polynomial time algorithms for the MIN CUT problem on degree restricted trees. *SIAM Journal on Computing*, 14(1):158–177, 1985.

[4] Michael R. Garey and David S. Johnson. *Computers and intractability. A guide to the theory of NP-completeness.* W. H. Freeman and Co., San Francisco, Calif., 1979.

[5] Martin Grohe, Ken-ichi Kawarabayashi, Dániel Marx, and Paul Wollan. Finding topological subgraphs is fixed-parameter tractable. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, (STOC 2011)*, pages 479–488, 2011.

[6] Dimitrios M. Thilikos Hans L. Bodlaender, Michael R. Fellows. Derivation of algorithms for cutwidth and related graph layout parameters. *Journal of Computer and System Sciences*, 75(4):231–244, 2009.

[7] Ephraim Korach and Nir Solel. Tree-width, path-width, and cutwidth. *Discrete Applied Mathematics*, 43(1):97–101, 1993.

[8] B. Monien and I. H. Sudborough. Min cut is NP-complete for edge weighted trees. *Theoretical Computer Science*, 58(1-3):209–229, 1988.

[9] Neil Robertson and Paul D. Seymour. Graph minors XXIII. Nash-Williams' immersion conjecture. *J. Combin. Theory Ser. B*, 100(2):181–205, 2010.

[10] Dimitrios M. Thilikos, Maria J. Serna, and Hans L. Bodlaender. Cutwidth I: A linear time fixed parameter algorithm. *Journal of Algorithms*, To appear.

[11] Mihalis Yannakakis. A polynomial algorithm for the min-cut linear arrangement of trees. *Journal of the ACM*, 32(4):950–988, 1985.