Geometric Extensions of Cutwidth in any Dimension

Menelaos Karavelas, Dimitris Zoros, Spyridon Maniatis, Dimitrios M. Thilikos

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 a hyperplane of $\mathbb{R}^d$ can intersect. Then, the d-cutwidth of G, denoted by $\text{CW}_d(G)$, is the minimun 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.

[§]AlGCo project team, CNRS, LIRMM, France.

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

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 $\mathbb{R}^d$ (which are known to be isomorphic to $\mathbb{R}^{d-1}$). Let Π be a hyperplane in $\mathbb{R}^d$, then there are $a_0, a_1, \dots, a_d \in \mathbb{R}$ such that $\Pi = \{(x_1, \dots, x_d) \in \mathbb{R}^d \mid a_1 x_1 + \cdots + a_d x_d + a_0 = 0\}$. 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 $\{(x_1, \dots, x_d) \in \mathbb{R}^d \mid \sum_{i=1}^{d} (x_i - c_i)^2 = r^2\}$. We denote by $S(d)$ the set of all hyperspheres of $\mathbb{R}^d$. We call a continuous function $C : [0, 1] \to \mathbb{R}^d$ a *curve* of $\mathbb{R}^d$ with ends $C(0)$ and $C(1)$.

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: (**a**) for every $e = \{u, v\} \in E$, the ends of C_e are $f(u)$ and $f(v)$, and (**b**) for all $x \in (0, 1)$ and $v \in V$ it holds that $f_e(x) \neq f(v)$. 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. 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 an element of $\mathbf{E}_d(G)$ *essential-embedding of G in $\mathbb{R}^d$*. Let $\mathcal{E}_d(G)$ be a 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).

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*

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

Observe that any hyperplane Π of $\mathbb{R}_d$ that meets a curve $C_e \in \mathcal{C}$ once, 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, for now on we will omit $\mathcal{C}$. Observe that the definition of 1-cutwidth, where hyperplanes degenerate to subspaces of $\mathbb{R}$ of dimension 0 (i.e., points) is equivalent to the usual definition of

[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.

cutwidth. Therefore, d-cutwidth is the intuitive generalization of the notion of cutwidth in any dimension $d \geq 2$. Also observe that our demand of essential embeddings is expressed here by our demand of injective functions.

3 Properties of d-cutwidth

This section is devoted to the last two statements of Theorem 1.

Proof of Theorem 1.iii. Consider the d-dimensional *moment curve* C with parametric equation $C(t) := (t, t^2, t^3, \ldots, t^d), \quad t \in \mathbb{R}$. Consider also 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

$$a_1x_1 + a_2x_2 + \ldots + a_dx_d + a_0 = 0, \quad \text{and} \quad x_i = t^i, \quad i = 1, \ldots, d$$

for t. This gives the polynomial equation $q(t) := a_0 + a_1t + a_2t^2 + \ldots + a_dt^d = 0$, 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., $\text{CW}_d(G) \leq d \ \text{CW}(G)$. $\square$

Spherical d-cutwith Given a graph $G = (V, E)$ and two positive integers k and d, where $d \geq 2$, we define the *spherical d-dimensional cutwidth* of G, or simply spherical d-cutwidth, to be equal to

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

The proof of Theorem 1.*iv* is a consequence of Theorem 1.*iii* and the following two lemmata.

Lemma 1 *For every graph G and any $d \geq 2$, $\text{CW}_d(G) \leq \text{SCW}_d(G) \leq (d+1)\ \text{CW}(G)$.*

Lemma 2 *For every graph G and every $d \geq 1$, $\text{CW}_{d+1}(G) \leq \text{SCW}_d(G)$.*

The above results clarify the relation between d-cutwidth and spherical d-cutwidth and we believe that they have independent interest. We omit the proofs as they are too lengthy to fit in this extended abstract.

4 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 $\text{CW}_d(G) \leq k$ 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 $\text{CW}_d(G) \leq k$ 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 $\text{CW}(G) \leq k$. If the answer is negative then we can safely report that $\text{CW}_d(G) > k$. 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.