Tree-width of hypergraphs and surface duality

Frédéric Mazoit

► To cite this version:

Frédéric Mazoit. Tree-width of hypergraphs and surface duality. DIMAP Workshop on Algorithmic Graph Theory, Mar 2009, France. pp.93-97, 10.1016/j.endm.2009.02.013 . hal-00380502

HAL Id: hal-00380502

<https://hal.science/hal-00380502>

Submitted on 1 May 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Tree-width of graphs and surface duality

Frédéric Mazoit^{1,2}

*LaBRI Université Bordeaux,
351 cours de la Libération F-33405 Talence cedex, France*

Abstract

In Graph Minors III, Robertson and Seymour conjecture that the tree-width of a graph and that of its dual differ by at most one. In this paper, we prove that given a hypergraph H on a surface of Euler genus k , the tree-width of H^* is at most the maximum of $\text{tw}(H) + 1 + k$ and the maximum size of a hyperedge of H^* minus one.

Keywords: Tree-width, duality, surface.

1 Introduction

Tree-width is a graph parameter introduced by Robertson and Seymour in connection with graph minors. In [RS84], they conjectured that for a planar graph G , $\text{tw}(G)$ and $\text{tw}(G^*)$ differ by at most one. In an unpublished paper, Lapoire [Lap96] proved a more general result: for any hypergraph H on an orientable surface of Euler genus k , $\text{tw}(H^*)$ is at most the maximum of $\text{tw}(H) + 1 + k$ and the maximum size of a hyperedge of H^* minus one. Nevertheless, his proof is rather long and technical. Later, Bouchitté et

¹ Email: Frederic.Mazoit@labri.fr

² Research supported by the french ANR-project "Graph decompositions and algorithms (GRAAL)".

al. [BMT03] gave an easier proof for planar graphs. Here we give an easy proof that Lapoire’s result holds for arbitrary surfaces.

2 Hypergraphs on surfaces and duality

A *surface* is a connected compact 2-manifold without boundaries. Oriented surfaces Σ can be obtained by adding “handles” to the sphere, and non-orientable surfaces, by adding “crosscaps” to the sphere. The *Euler genus* or just *genus* $k(\Sigma)$ of Σ is twice the number of handles added if Σ is orientable and $k(\Sigma)$ is the number of crosscaps added otherwise. We denote by $\bar{X}$ the closure of a subset X of Σ . Two disjoint subsets X and Y of Σ are *incident* if $X \cap \bar{Y}$ or $Y \cap \bar{X}$ is non empty.

A graph on a surface Σ is a drawing of an abstract graph on Σ , i.e. each vertex is an element of Σ , each edge is an open curve between two vertices, and edges are pairwise disjoint. A bipartite graph $G = (V \cup V_E, L)$ on Σ can be seen as the incidence graph of a hypergraph. For each $v_e \in V_E$, we merge v_e and its incident edges into a *hyperedge* e and call v_e its *center*. Let E be the set of all *hyperedges*. A *hypergraph on Σ* is any such pair $H = (V, E)$. We often contract hyperedges in *edges*, and we only consider graphs and hypergraphs up to homeomorphism.

A *face* of a hypergraph H on Σ is a connected component of $\Sigma \setminus H$. We denote by $V(H)$, $E(H)$ and $F(H)$ the vertex, edge and face sets of H . The elements of $A(H) = V(H) \cup E(H) \cup F(H)$ are the *atoms* of H , they partition Σ . We also consider graphs and hypergraphs on surfaces as abstract graphs or hypergraphs. For example, we consider an edge e as a subset of Σ or as a set of vertices. The maximum size of an edge of H is $\alpha(H)$. A *cut-edge* in a hypergraph H on Σ is an edge e that “separates” H , i.e. H intersects at least two connected components of $\Sigma \setminus \bar{e}$. As an example, if a planar graph G has a cut-vertex u , any loop on u that goes “around” a connected component of $G \setminus \{u\}$ is a cut-edge. In the following, we only consider *2-cell* hypergraphs, i.e. hypergraphs whose faces are homeomorphic to open discs. Euler’s formula links the number of vertices, edges and faces of a 2-cell graph G to $k(\Sigma)$:

$$|V(G)| - |E(G)| + |F(G)| = 2 - k(\Sigma).$$

The dual of a hypergraph $H = (V, E)$ on Σ is obtained by choosing a vertex v_f in every face f of H , and for every edge e of center v_e , we pick up an edge e^* as follows. Choose a local orientation of the surface around v_e . This local orientation induces a cyclic order $v_1, f_1, v_2, f_2, \dots, v_d, f_d$ of the ends of e and

of the faces incident with e (possibly with repetition). The edge e^* is the edge obtained by “rotating” e and whose ends are $v_{f_1}, \dots, v_{f_d}$.

In the following, we suppose, for simplicity, that H has no cut-edge.

3 P-trees and tree-decompositions

A *tree-decomposition* of a hypergraph H is a pair $\mathcal{T} = (T, (X_v)_{v \in V(T)})$ with T a tree and $(X_v)_{v \in V(T)}$ a family of subsets of vertices of H called *bags* with:

- i. $\bigcup_{v \in V(T)} X_v = V(H)$;
- ii. $\forall e \in E(H), \exists v \in V(T)$ with $e \subseteq X_v$;
- iii. $\forall x, y, z \in V(T)$ with y on the path from x to z , $X_x \cap X_z \subseteq X_y$.

The *width* of $\mathcal{T}$ is $\text{tw}(\mathcal{T}) = \max(|X_t| - 1; t \in V(T))$ and the *tree-width* $\text{tw}(H)$ of H is the minimum width of one of its tree-decompositions.

The *border* of a partition μ of E is the set of vertices $\delta(\mu)$ that are incident with edges in at least two parts of μ , and the border of $E' \subseteq E$ is the border of $\{E', E \setminus E'\}$. A partition $\mu = \{E_1, \dots, E_p\}$ of E is *connected* if there is a partition $\{V_1, E_1, F_1, \dots, V_p, E_p, F_p\}$ of $A(H) \setminus \delta(\mu)$ so that each $V_i \cup E_i \cup F_i$ is connected in Σ . A *labelled tree* of H is a tree T whose leaves are labelled by edges of H in a bijective way. Removing an internal node v of T results in a partition of the leaves of T and thus in a *node-partition* λ_v of E . Keeping the leaf labels of T and labelling each internal node v of T with $\delta(\lambda_v)$ turns T into a tree-decomposition. The *tree-width* of a labelled tree is its *tree-width*, seen as a tree-decomposition. A *p-tree* is a labelled tree whose internal vertices have degree three and whose node partitions are connected.

Let $\{A, B\}$ be a connected bipartition of H and $\{V_A, A, F_A, V_B, B, F_B\}$ a corresponding partition of $A(H) \setminus \delta(\{A, B\})$. We define H/A , the hypergraph H in which the edges in A are contracted into a new edge $e_A = \delta(\{A, B\})$ by mean of its incidence graph as follows. Let $G_H = (V \cup V_E, L)$ be the incidence graph of H . Identify the edges in A with their centers. By adding edges through faces in F_A , we can make $G_H[A \cup V_A]$ connected. We then contract $A \cup V_A$ into a single edge center v_A . To make the resulting graph bipartite, we remove all v_A -loops. When removing a loop e , the merged face may not be a disc. In this case, we “cut” Σ along the border of the merged face and fill the holes with discs. This operation decreases the genus of the surface. A connected partition $\{A, B\}$ is *trivial* if H/A or H/B is equal to H .

4 The main theorem

We need the following folklore lemma to prove Proposition 4.2.

Lemma 4.1 *For any connected bipartition $\{A, B\}$ of a hypergraph H on a surface Σ , $\text{tw}(H) \leq \max(\text{tw}(H/A), \text{tw}(H/B))$. If $\delta(\{A, B\})$ belongs to a bag of an optimal tree-decomposition, then $\text{tw}(H) = \max(\text{tw}(H/A), \text{tw}(H/B))$.*

Proposition 4.2 *There exists a p-tree T of H with $\text{tw}(T) = \text{tw}(H)$.*

Proof. By induction on $|E|$, if $|E| \leq 3$, the only labelled tree is an optimal p-tree. Otherwise, we claim that there exists a connected non trivial bipartition $\{A, B\}$ of E whose border is contained in a bag of an optimal tree-decomposition of H . Since $\{A, B\}$ is connected, neither e_A nor e_B are cut-edges in H/A and H/B . By induction, there exist p-trees T_A and T_B of H/A and H/B , each of optimal width. By removing the leaves labelled e_A and e_B and adding an edge between their respective neighbours, we obtain from $T_A \sqcup T_B$ a p-tree whose width is $\max(\text{tw}(T/A), \text{tw}(T/B))$ which is equal, by Lemma 4.1, to $\text{tw}(H)$. $\square$

Because of the natural bijection between $E(H)$ and $E(H^*)$, a p-tree T of H also corresponds to a labelled tree T^* of H^* .

Proposition 4.3 *For any p-tree T of H ,*

$$\text{tw}(T^*) \leq \max(\text{tw}(T) + 1 + k(\Sigma), \alpha(H^*) - 1).$$

Proof. Let v be a vertex of T labelled X_v in T and X_v^* in T^* . If v is a leaf, then $X_v^* = e^*$ and $|X_v^*| - 1 = |e^*| - 1 \leq \max(\text{tw}(T) + 1 + k(\Sigma), \alpha(H^*) - 1)$.

We can suppose that v is an internal node whose partition is $\{A, B, C\}$. The labels of v in T and T^* are respectively $X_v = \delta(\{A, B, C\})$ and X_v^* , the set of faces incident with edges in at least two parts among A , B and C . We want to transform the incidence graph G_H of H to remove all vertices in $V(H) \setminus X_v$ and all the faces that do not belong to X_v^* . To do so, we proceed as for the proof of Proposition 4.2, by contracting A (and B and C) which is possible because $\{A, B, C\}$ is connected. But since we now care about the faces in X_v^* , we have to be more careful. We may add faces to X_v^* but not remove faces from it, so we can add edges to make say $G_H[A \cup V_A]$ connected, but we can not remove a loop e on v_A incident with two faces in X_v^* . To remove such a loop e , we cut Σ along e and fill the holes with open discs that we can contract. During this process, we cut v_A in two *siblings*, and decrease the genus of Σ .

After contracting A , B and C , we obtain a bipartite graph G_v on a surface Σ' that has $|X_v| + 3 + s$ vertices with s the number of siblings, at least $|X_v^*|$ faces and

with $k(\Sigma') \leq k(\Sigma) - s$. Since G_v is bipartite and faces in X_v^* are incident with at least 4 edges, $2|E(G_v)| = 4|F_4| + 6|F_6| + \dots \geq 4|F(G_v)|$ with F_{2k} the set of $2k$ -gon faces of G_v , and thus $|E(G_v)| \geq 2|F(G_v)|$. If we apply Euler's formula to G_v , we obtain: $|X_v| + 3 + s - |E(G_v)| + |F(G_v)| = 2 - k(\Sigma') \geq 2 - k(\Sigma) + s$. Adding this to $|E(G_v)| \geq 2|F(G_v)|$, we get $|X_v| + 1 + k(\Sigma) \geq |F(G_v)| \geq |X_v^*|$ which proves that $|X_v^*| - 1 \leq \max(\text{tw}(T) + 1 + k(\Sigma), \alpha(H^*) - 1)$, and thus $\text{tw}(T^*) \leq \max(\text{tw}(T) + 1 + k(\Sigma), \alpha(H^*) - 1)$. $\square$

Our main theorem is a direct corollary of Proposition 4.2 and Proposition 4.3.

Theorem 4.4 *For any hypergraph H on a surface Σ ,*

$$\text{tw}(H^*) \leq \max(\text{tw}(H) + 1 + k(\Sigma), \alpha(H^*) - 1).$$

5 Conclusion

A graph on a surface is not likely to be much more complicated than its dual. Our theorem shows that, for a graph G on a surface Σ , $|\text{tw}(G^*) - \text{tw}(G)| \leq 1 + k(\Sigma)$, and thus that $\text{tw}(G)$ and $\text{tw}(G^*)$ are roughly the same, which shows that tree-width is indeed quite a robust complexity parameter for graphs.

In Graph Minors III, Robertson and Seymour gave an example of dual graphs whose respective tree-widths differ by one, and thus meet the given bound. We have not been able to find such pairs for higher genus which raises the question of the optimality of this bound. We conjecture that it is not optimal.

References

- [BMT03] V. Bouchitté, F. Mazoit, and I. Todinca. Chordal embeddings of planar graphs. *Discrete Mathematics*, 273:85–102, 2003.
- [Lap96] D. Lapoire. Treewidth and duality for planar hypergraphs. Manuscript http://www.labri.fr/perso/lapoire/papers/dual_planar_treewidth.ps, 1996.
- [RS84] N. Robertson and P. D. Seymour. Graph Minors. III. Planar Tree-Width. *Journal of Combinatorial Theory Series B*, 36(1):49–64, 1984.