



HAL
open science

Tree-width of hypergraphs and surface duality

Frédéric Mazoit

► **To cite this version:**

Frédéric Mazoit. Tree-width of hypergraphs and surface duality. *Journal of Combinatorial Theory, Series B*, 2010, 102 (3), pp.671-687. 10.1016/j.jctb.2011.11.002 . hal-00492498v2

HAL Id: hal-00492498

<https://hal.science/hal-00492498v2>

Submitted on 18 Nov 2011

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 hypergraphs and surface duality¹

Frédéric Mazoit²

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

Abstract

In Graph Minors III, Robertson and Seymour write: “It seems that the tree-width of a planar graph and the tree-width of its geometric dual are approximately equal — indeed, we have convinced ourselves that they differ by at most one.” They never gave a proof of this. In this paper, we prove a generalisation of this statement to embedding of hypergraphs on general surfaces, and we prove that our bound is tight.

Keywords: tree-width, duality, surface.

1 Introduction

Tree-width is a graph parameter which was first defined by Halin [Hal76], and which has been rediscovered many times (see [AP89, RS84]). In [AP89], Arnborg and Proskurovski introduced a general framework to efficiently solve NP-complete problems when restricted to graphs of bounded tree-width. Courcelle [Cou90] extended this framework by showing that any problem expressible in a certain logic can be efficiently solved for graphs of bounded tree-width. Tree-width thus seems to be a good “complexity measure” for graphs.

Given an embedding Γ of a graph in a surface, it is easy to obtain the dual embedding Γ^* : just put a vertex in each face and for every edge e separating the faces f and g , add a dual edge fg . One could thus expect that Γ and Γ^* have the same “complexity”, and indeed in [RS84], Robertson and Seymour claimed that for a plane embedding Γ , $\text{tw}(\Gamma)$ and $\text{tw}(\Gamma^*)$ differ by at most one.

In an unpublished paper, Lapoire [Lap96] gave a more general statement about embeddings of hypergraphs on orientable surfaces. Nevertheless,

¹This research is supported by the french ANR project DORSO.

²Email: Frederic.Mazoit@labri.fr

his proof was rather long and technical. Later, Bouchitté et al. and Mazoit [BMT03, Maz04] gave easier proofs for plane graphs. Here we give a proof that Lapoire’s claim is valid for general surfaces³:

Theorem 1 *For any 2-cell embedding Λ of a hypergraph on a surface Σ ,*

$$\text{tw}(\Lambda^*) \leq \max\{\text{tw}(\Lambda) + 1 + k_\Sigma, \alpha_{\Lambda^*} - 1\}.$$

Here α_{Λ^*} is the maximum size of an edge of Λ^* and k_Σ is the Euler genus of Σ .

In Section 2, we give the basic definitions. Section 3 is devoted to the proof of Theorem 1 while in Section 4 we give examples of embeddings which match the bound of this theorem.

2 Preliminaries

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* such that every edge of H is contained in at least one bag of \mathcal{T} , and for every vertex $v \in V_H$, the vertices of T whose bag contain v induces a non-empty sub-tree of T . 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.

A *surface* is a connected compact 2-manifold without boundaries. Oriented surfaces can be obtained by adding “handles” to the sphere, and non-orientable surfaces, by adding “crosscaps” to the sphere. The *Euler genus* k_Σ of a surface Σ (or just *genus*) is twice the number of handles if Σ is orientable, and is the number of crosscaps otherwise.

We denote by \overline{X} the closure of a subset X of Σ . We say that two disjoint subsets X and Y of Σ are *incident* if $X \cap \overline{Y}$ or $Y \cap \overline{X}$ is non-empty. Since we consider finite graphs and hypergraphs, we can assume that the curves involved in the embeddings are not completely wild and are, say, piecewise linear. This implies that connectivity and arc-connectivity coincide. An *open curve* is a subset of Σ which is homeomorphic to $]0, 1[$. An open curve whose closure is homeomorphic to the 1-sphere S^1 is a *loop* and is a *strait edge* otherwise. A connected subset X of Σ is a *star* if it contains a point v_X called its *centre* such that $X \setminus \{v_X\}$ is a union of pairwise disjoint strait edges called *half edges*. Note that an open curve is also a star. Let X be an open curve or a star. The elements of $\overline{X} \setminus X$ are the *ends* of X .

³This result also appears as an extended abstract in [Maz09]. Unfortunately although the general scheme of the proof is the same, some definitions are wrong and we could not obtain a valid proof with them.

An *embedding of a hypergraph on a surface* Σ is a pair $\Lambda = (V, E)$ in which V is a finite subset of Σ whose elements are the *vertices* of the embedding, and E is a finite set of pairwise disjoint stars called (*hyper*)*edges*. Edges contain no vertex and their ends are vertices. Such an embedding naturally corresponds to an abstract hypergraph H . We say that Λ is an *embedding of H* . An *embedding of a graph on Σ* is an embedding of a hypergraph whose edges are straight edges and loops. Let Λ be an embedding of a hypergraph on Σ . We denote by V_Λ the vertex set of Λ and by E_Λ the edge set of Λ . Let V_{E_Λ} contain the centre of all the edges and let L_Λ contain all the half edges. Then $(V_\Lambda \cup V_{E_\Lambda}, L_\Lambda)$ is an embedding of a bipartite graph on Σ which is the *incidence embedding of Λ* . We denote embeddings of graphs with the Greek letters Γ and Π and embeddings of hypergraphs with the Greek letter Λ . We only consider embedding of graphs and hypergraphs up to homeomorphisms. Since embeddings of hypergraphs naturally have abstract counterparts, we apply graph theoretic notions to them without further notice. For example, we may consider an edge e as a subset of Σ or as a set of vertices. We also consider embeddings of hypergraphs on Σ as subsets of Σ .

A *face* of an embedding Λ is a component of $\Sigma \setminus \Lambda$. We denote by F_Λ the set of faces of Λ . An embedded hypergraph is *2-cell* if all its faces are homeomorphic to open discs. Let Γ be a 2-cell embedding of a graph on a surface Σ . Euler's formula links the number of vertices, edges and faces of Γ and the genus of the surface:

$$|V_\Gamma| - |E_\Gamma| + |F_\Gamma| = 2 - k_\Sigma.$$

We now let Λ be a 2-cell embedding of a hypergraph on a surface Σ . The *dual* of Λ is the embedding Λ^* such that:




Figure 1: A planar hypergraph and its dual.

- i. Every vertex of Λ^* belongs to a face of Λ and every face of Λ contains exactly one vertex of Λ^* ;

- ii. For every edge e of Λ , there exists a dual edge e^* sharing its centre, and every edge of Λ^* corresponds to an edge of Λ .
- iii. For every edge e of Λ with centre v_e , the half edges of e and e^* around v_e alternate in their cyclic order.

Note that the construction does not need Λ to be 2-cell but if not, Λ^* is not unique and $(\Lambda^*)^*$ need not be Λ .

3 The upper bound

Since Theorem 1 is about 2-cell embeddings, and since the theorem is trivial for edge-less embeddings, we always consider connected embeddings and hypergraphs with at least one edge.

The *border* of a partition μ of E_H is the set $\delta_H(\mu)$ of vertices which are incident with edges in at least two parts of μ , and the *border* of $A \subseteq E_H$ is $\delta_H(A) = \delta_H(\{A, E_H \setminus A\})$. A *partitioning tree* of H is a tree T whose leaves are bijectively labelled by edges of H . 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_H . Removing an edge e of T results in a bipartition of the leaves of T and thus in an *edge-partition* λ_e of E_H .

Lemma 1 *Let H be a connected hypergraph with at least one edge. Let T be a partitioning tree of H . Labelling each internal node v of T with $\delta_H(\lambda_v)$ turns T into a tree-decomposition.*

Proof. By construction, every edge of H is contained in a bag of T . Let $x \in V_H$. Let S be the set of leaves of T whose label contain x , and let T_x be the subtree of T whose set of leaves is S . Since x is not isolated, T_x contains at least one leaf. Moreover, an internal bag of T contains x if and only if it separates two leaves u and v of T whose edge label contain x . Since those bags are precisely its internal bags, T_x is precisely the subgraph induced by the vertices of T whose bag contain x . \square

The *tree-width* of a partitioning tree is its *tree-width*, seen as a tree-decomposition.

Let Λ be a 2-cell embedding of a hypergraph on a surface Σ . If T is a partitioning tree of Λ , then *dual* of T is the partitioning tree T^* of Λ^* obtained by replacing in T each label e by the dual edge e^* .

Given these definition, it is tempting to try to prove that for any embedding Λ of a hypergraph on Σ :

- i. there always exists a partitioning tree T such that $\text{tw}(T) = \text{tw}(\Lambda)$;
- ii. for any partitioning tree T , $\text{tw}(T^*) \leq \max\{\text{tw}(T) + 1 + k_\Sigma, \alpha_{\Lambda^*} - 1\}$.

The first item is true but we could not prove the second one. However, we prove that both properties hold for a restricted class of partitioning tree which we call p-trees.

3.1 A sketch of the planar case

Before we go on with the proof, we consider the planar case as it contains most ideas. The proof which we now sketch is based on the proof in [Maz04]. Note that all definitions in this subsection are local to this subsection.

Let Γ be an embedding of a graph on the sphere S^2 . Moreover, let us suppose that Γ has no bridge and no loop. A *pretty curve* is a subset of S^2 which is homeomorphic to S^1 , which crosses Γ only on vertices, and which never “enters” a face twice or more. A Θ -*structure* is a union of three curves $\rho_e \cup \rho_f \cup \rho_g$ such that $\rho_e \cup \rho_f$, $\rho_f \cup \rho_g$ and $\rho_g \cup \rho_e$ are all pretty curves. Pretty curves induce bipartitions of E_Γ , and Θ -structures induce tripartitions of E_Γ . A partitioning tree is *geometric* if all its node partitions come from Θ -structures.

Let T and T^* be dual geometric partitioning tree, and let v be a node of T . We claim that the size of the dual bags X_v and X_v^* differ by at most 1. This is clearly true for leaf bags whose size is either 1 or 2. So let us suppose that v is an internal node and let $\Theta = \rho_e \cup \rho_f \cup \rho_g$ be a Θ -structure realising the node partition λ_v . By construction X_v contains all the vertices which belong to Θ , and X_v^* contains all the faces that Θ goes through. Since $\rho_e \cup \rho_f$ is a pretty curve which alternatively crosses vertices and faces of Γ and never enters the same face twice, $|X_v \cap (\rho_e \cup \rho_f)| = |X_v^* \cap (\rho_e \cup \rho_f)|$. The difference between $|X_v|$ and $|X_v^*|$ thus comes from ρ_g . But ρ_g also alternatively crosses vertices and faces of Γ without entering the same face twice. This implies that difference between $|X_v|$ and $|X_v^*|$ is at most 1. Since this inequality holds for any node, we have $\text{tw}(T^*) \leq \text{tw}(T) + 1$.

To finish, we only need to prove that there exists a geometric partitioning tree T such that $\text{tw}(\Gamma) = \text{tw}(T)$. To do this, we apply an induction on planar hypergraphs. Suppose that ρ is a pretty curve whose bipartition of E_Γ is $\{A, B\}$. Let D_A and D_B be the two components of $S^2 \setminus \rho$. If we remove all the vertices and edges in D_A and replace them by a star whose set of ends is $\delta_\Gamma(A)$, we obtain a contracted hypergraph $\Gamma_{/A}$. Let $T_{/A}$ and $T_{/B}$ be geometric partitioning trees of $\Gamma_{/A}$ and $\Gamma_{/B}$. By removing from the disjoint union $T_{/A} \cup T_{/B}$ the leaves labelled e_A and e_B and adding a new edge between their

respective neighbours, we obtain a geometric partitioning tree T of Γ . We show that $\text{tw}(T) = \max\{\text{tw}(T_{/A}), \text{tw}(T_{/B})\}$. By induction $\text{tw}(\Gamma_{/A}) = \text{tw}(T_{/A})$ and $\text{tw}(\Gamma_{/B}) = \text{tw}(T_{/B})$. The result follows from the fact that it is always possible to find ρ such that $\text{tw}(\Gamma) = \max\{\text{tw}(\Gamma_{/A}), \text{tw}(\Gamma_{/B})\}$.

Before we go on with the general case, let us make some comments.

- On higher genus surfaces, we can still describe separators in terms of curves on Σ but as the genus increases, the number of curves involved increases and it quickly becomes too complex to control how the curves interact. As a matter of fact we do not really care about curves. Θ -structures are only important because they cut the sphere in three connected regions.

Indeed, we can prove that $|X_v^*| \leq |X_v| + 1$ without considering curves as follows. Let Θ be a Θ -structure which realises λ_v , and let D_A , D_B and D_C be the components of $S^2 \setminus \Theta$. We contract all the vertices and faces which are contained in D_A into a single vertex v_A . We do the same thing in D_B and D_C and to obtain two vertices v_B and v_C . We obtain a bipartite embedding Γ_v whose set of vertices is $\{v_A, v_B, v_C\} \cup X_v$ and whose set of faces is X_v^* . Euler's formula thus implies that $(|X_v| + 3) + |X_v^*| - |E_{\Gamma_v}| = 2$. Since the faces of Γ_v are incident with at least 4 vertices, it is easy to prove that $|E_{\Gamma_v}| \geq 2|X_v^*|$. The bound for the planar case follows.

Although the contracting process becomes quite technical, this proof does generalise to higher genus surfaces.

- Some loops and bridges are troublesome and must be taken care of separately.

Indeed, let e be a loop of Γ which separates vertices of Γ , and let v be an internal node of T which is the neighbour of a leaf labelled by e . Since any curve which isolates e from $E \setminus \{e\}$ has to enter the end of e twice, the node partition λ_v cannot come from a Θ -structure. The same kind of problem arises if e is a separating bridge because any curve isolating e must enter the same face twice.

Let e be a separating loop. If we take a closer look, the bag X_v is $\{e\}$ so the bag X_v^* should be $\{e^*\}$. For such internal nodes, we could just drop the condition λ_v comes from a Θ -structure and take any partition whose border is contained in e . But this idea does not work. For example, in figure 2, e is a separating loop whose end is v . The border of the partition $\{\{a\}, \{b, c, d\}, \{e\}, \{f\}, \}$ is $\{v\}$ but the border of the dual partition is $\{F1, F2, F3\}$ whereas the dual of e is the edge $\{F1, F2\}$.




Figure 2: A troublesome loop.

To make this work, we cannot just take any partition whose border is a subset of e , we have to take a partition whose parts correspond to the connected components of $S^2 \setminus \bar{e}$.

- In the planar case, for our purpose, any two pretty curve which induce the same bipartition of E_Γ are equivalent. To avoid explicitly considering equivalent classes of pretty curves, the proof in [Maz04] proceed as follows. For each face $F \in F_\Gamma$, it put a vertex which is linked to all the vertices in V_Γ which are incident to F . The resulting embedding Π a *radial embedding* of Γ . Pretty curves then correspond to cycles of Π .

In this paper, we do not explicitly realise partitions of E_Γ with curves on Σ but with some disjoint connected subsets $\Sigma_A, \Sigma_B, \Sigma_C$ of Σ . As for curves, there is not a single way to realise a node partition with such subsets and we use radial embeddings to avoid dealing with cumbersome equivalent classes.

3.2 Partitioning trees

Given a non-empty subset $A \subseteq E_H$, we define the *contracted hypergraph* $H_{/A}$ of H as the hypergraph with vertex set $\cup(E_H \setminus A)$ and with edge set $(E_H \setminus A) \cup \{e_A\}$ in which $e_A = \delta_H(A)$ is a new hyperedge. Let $\{A, B\}$ be a non trivial bipartition of E_H and $T_{/A}$ and $T_{/B}$ be respectively partitioning trees of $H_{/A}$ and $H_{/B}$. By removing from the disjoint union $T_{/A} \cup T_{/B}$ the leaves labelled e_A and e_B and adding a new edge between their respective neighbours, we obtain a partitioning tree T which is the *merge of $T_{/A}$ and $T_{/B}$* .

Lemma 2 *Let H be a connected hypergraph with at least one edge. Let $\{A, B\}$ be a non trivial bipartition of E_H , and let $T_{/A}$ and $T_{/B}$ be partitioning trees of $H_{/A}$ and $H_{/B}$. Then the merge T of $T_{/A}$ and $T_{/B}$ is such that*

$$\text{tw}(T) = \max\{\text{tw}(T_{/A}), \text{tw}(T_{/B})\}.$$

Proof. Let $C \subseteq E_H$ be disjoint from A . We claim that $\delta_H(C)$ and $\delta_{H_{/A}}(C)$ are equal. Indeed, let $v \in \delta_{H_{/A}}(C)$. By definition, there exists $e \in E_{H_{/A}} \setminus C$

and $f \in C$ which contain v . If $e \neq e_A$, then $e \in E_H \setminus C$. Otherwise, $e = e_A = \delta_H(A)$ and there exists $e' \in A \subseteq E_H \setminus C$ which contains v . In both cases, $v \in \delta_H(C)$. Conversely, let $v \in \delta_H(C)$. By definition, there exists $e \in E_H \setminus C$ and $f \in C$ which contain v . If $e \notin A$, then $e \in E_{H/A} \setminus C$. Otherwise $v \in \delta_H(A) = e_A$. In both cases $v \in \delta_{H/A}(C)$.

Let u be an internal node of T . By symmetry, we can suppose that u belongs to T/A . The node-partition of u in T/A is $\lambda_{u/A} = \{E_1 \cup \{e_A\}, E_2, \dots, E_p\}$, and the node-partition of u in T is $\lambda_u = \{E_1 \cup A, E_2, \dots, E_p\}$. The above claim implies that $\delta_H(\lambda_u) = \delta_{H/A}(\lambda_{u/A})$. The result follows. \square

Lemma 2 and the following folklore lemma are the key tools to our proof of Theorem 2 that there always exists a p-tree of optimal width.

Lemma 3 *Let H be a hypergraph with at least one edge and no isolated vertices. For any bipartition $\{A, B\}$ of E_H ,*

$$\text{tw}(H) \leq \max\{\text{tw}(H/A), \text{tw}(H/B)\}.$$

If $\delta_H(\{A, B\})$ belongs to a bag of an optimal tree-decomposition of H , then

$$\text{tw}(H) = \max\{\text{tw}(H/A), \text{tw}(H/B)\}.$$

Proof. Let $\mathcal{T}/A = (T/A, (X_v)_{v \in V_{T/A}})$ and $\mathcal{T}/B = (T/B, (Y_v)_{v \in V_{T/B}})$ be respective optimal tree-decompositions of H/A and H/B . Let u be a vertex of T/A whose bag contain e/A and let v be a vertex of T/B whose bag contain e/B . By adding an edge uv to the disjoint union $T/A \cup T/B$, we obtain a tree-decomposition \mathcal{T} of H such that $\text{tw}(\mathcal{T}) = \max\{\text{tw}(H/A), \text{tw}(H/B)\}$, which proves the first part of the lemma.

Suppose now that $\delta_H(\{A, B\})$ belongs to the bag of a vertex v of an optimal tree-decomposition $\mathcal{T} = (T, (Z_v)_{v \in V_T})$ of H . By removing $V \setminus (\cup B)$ from the bags of \mathcal{T} , we obtain a tree-decomposition \mathcal{T}/A of H/A such that $\text{tw}(\mathcal{T}/A) \leq \text{tw}(H)$. Similarly, we obtain a tree-decomposition \mathcal{T}/B of H/B such that $\text{tw}(\mathcal{T}/B) \leq \text{tw}(H)$. The second part of the lemma follows. \square

3.3 Radial embeddings

Note that, in this subsection, we do not require embedded hypergraphs to be 2-cell but they must be connected and have at least one edge.

Let Λ be an embedding of a hypergraph on a surface Σ . A *radial embedding* of Λ is an embedding Π of a bipartite graph on Σ such that:

- i. $\{V_\Lambda, V_\Pi \setminus V_\Lambda\}$ is a bipartition of V_Π , and $V_\Pi \setminus V_\Lambda$ contains exactly one vertex per face of Λ ;




Figure 3: A radial embedding of the left example of Figure 1.

- ii. each edge of Λ is contained in a face of Π and each face of Π contains exactly one edge of Λ .

First radial embeddings do exist.

Lemma 4 *Every embedding Λ of a connected hypergraph with at least one edge on a surface Σ admits a radial embedding.*

Proof. The set V_Λ being fixed, let us first choose one *face vertex* per face of Λ to get $V_\Pi \setminus V_\Lambda$. Let $(D_e)_{e \in E_\Lambda}$ be pairwise disjoint open discs such that each D_e contains e . Such discs can be obtained by “thickening” each edge a little. We now continuously distort all the discs intersecting a given face so that they become incident with its corresponding “face vertex”.

The union of the boundaries of the discs D_e correspond to the drawing of a bipartite graph Γ that satisfies all the required condition except that some faces may be empty. Indeed, suppose that $F \in F_\Lambda$ is homeomorphic to a disc and that $v \in V_\Lambda$ is incident with F . Let e_1 and $e_2 \in E_\Lambda$ bound an “angle at v in F ”. Between e_1 and e_2 , there is an edge $f_1 \in E_\Gamma$ which is in the boundary of D_{e_1} and and edge $f_2 \in E_\Gamma$ which is in the boundary of D_{e_2} . The edges f_1 and f_2 bound an empty face.

As long as Γ contains an empty edge F , we remove any edge incident with F to merge it with a neighbouring face and thus decrease the total number of empty faces of Γ . In the end, we obtain a radial embedding Π of Λ . \square

If Λ is a 2-cell embedding, then the radial embedding of Λ is unique and two distinct embeddings share the same radial embedding if and only if they are dual embeddings. But, as already mentioned, we consider embeddings which are not 2-cell. This implies that a given embedding may have more than one radial embedding (see Figure 4).

Let Λ be an embedding of a hypergraph on a surface Σ , and let Π be a radial embedding of Λ . We say that an edge or a vertex of Π is *private to* a set F of faces of Π if all the faces it is incident to belong to F .




Figure 4: Two radial embeddings of a hypergraph on the projective plane.

We now define several notions *with respect to a radial embedding* Π . Let A be a set of edges of Λ . We denote by A_Π the open set that contains all the faces of Π corresponding to edges in A together with the edges and vertices of Π which are private to these faces. We say that A is Π -connected if A_Π is connected, and that a partition of E is Π -connected if its parts are Π -connected. Two edges e and f of Λ are Π -adjacent if $\{e, f\}_\Pi$ is Π -connected. An edge e of Λ is *troublesome* if the partition $\{e, E \setminus \{e\}\}$ is not Π -connected. The components of $\{e\}_\Pi$ then induce a partition of $E_\Lambda \setminus \{e\}$. Together with $\{e\}$, this partition is the *e-partition*.

If a vertex of Π is private to a set of faces of Π , then so are all its incident edges. Thus if we denote by G^Π the graph whose vertices are the edges of Λ , and in which two vertices are adjacent if their corresponding edges are Π -adjacent, then Π -connected sets of edges of Λ exactly correspond to connected subgraphs of G^Π .

Let A be a Π -connected set of edges of Λ . Let us denote \tilde{A} and H the respective abstract counterparts of A and Λ . If we remove the edges and vertices of Λ which are contained in A_Π and replace them by an edge e_A whose set of ends is $\delta_\Lambda(A)$ (which is possible because A_Π is connected), we obtain an embedding of a hypergraph whose abstract counterpart is $H_{/\tilde{A}}$. We thus denote this new embedding by $\Lambda_{/A}$. By removing from Π all the edges and vertices which are contained in A_Π , we obtain the *contracted radial embedding* $\Pi_{/A}$ of $\Lambda_{/A}$.

Remark 1 Let $A \subseteq E_\Lambda$ be Π -connected. By construction of $\Pi_{/A}$,

- i. a partition $\{E_1 \cup A, E_2, \dots, E_l\}$ is Π -connected if and only if the partition $\{E_1 \cup \{e_A\}, E_2, \dots, E_l\}$ is $\Pi_{/A}$ -connected;

- ii. edge e is troublesome in Λ_A if and only if e is a troublesome edge in Λ (and thus e belongs to Λ);
- iii. moreover, $\{\{e\}, E_2 \cup \{e_A\}, E_3, \dots, E_p\}$ is the e -partition in $\Lambda_{/A}$ if and only if $\{\{e\}, E_2 \cup A, E_3, \dots, E_p\}$ is the e -partition in Λ .

3.4 P-trees

Let Λ be an embedding of a connected hypergraph with at least one edge on a surface Σ , and let Π be a radial embedding of Λ . A p -tree of (Λ, Π) is a partitioning tree T of Λ such that:

- i. if a edge-partition λ_e of T is not Π -connected, then e is incident with a leaf labelled by a troublesome edge.
- ii. if v is an internal node of T whose degree is not 3, then v is a neighbour of a leaf labelled by a troublesome edge e and λ_v is the e -partition.

Note that when no troublesome edge exist, all edge partitions are Π -connected and all internal nodes have degree three.

Remark 1 implies that:

Lemma 5 *Let Λ be an embedding of a connected hypergraph with at least one edge on a surface Σ , let Π be a radial embedding of Λ , and let $\{A, B\}$ be a Π -connected bipartition of E_Λ . Let $T_{/A}$ and $T_{/B}$ be p -tree of $(\Lambda_{/A}, \Pi_{/A})$ and $(\Lambda_{/B}, \Pi_{/B})$ respectively. The merged partitioning-tree T of $T_{/A}$ and $T_{/B}$ is a p -tree of (Λ, Π) .*

Proof. Let e be an edge of T , and let λ_e be its edge-partition in T . If e is the edge linking $T_{/A}$ and $T_{/B}$, then $\lambda_e = \{A, B\}$ is Π -connected. Otherwise, by symmetry, we can suppose that, say, e belong to $T_{/A}$. By Remark 1.i., if edge-partition of e in $\Lambda_{/A}$ is $\Pi_{/A}$ -connected, then the edge-partition of e in Λ is Π -connected. Otherwise, say, e is incident with a leaf labelled by a troublesome edge in $\Lambda_{/A}$, and by Remark 1.ii., e is incident with a leaf labelled by a troublesome edge in Λ . Moreover, Remark 1.iii. directly implies that the node condition of p -trees is satisfied. \square

We can now prove that

Theorem 2 *Let Λ be an embedding of a connected hypergraph with at least one edge on a surface Σ , and let Π be a radial embedding for Λ . There exists a p -tree T of (Λ, Π) such that $\text{tw}(T) = \text{tw}(\Lambda)$.*

Proof. We proceed by induction on $|V_\Lambda| + |E_\Lambda|$. We say that an embedding Λ is *smaller* than another embedding Λ' if $|V_\Lambda| + |E_\Lambda| < |V_{\Lambda'}| + |E_{\Lambda'}|$. Let \mathcal{T} be a tree-decomposition of Λ of optimal width. We may assume that T has no two neighbouring bags, one of which contains the other — otherwise, contract the corresponding edge in T and merge the bags.

Suppose that we find a Π -connected bipartition $\{A, B\}$ of E_Λ whose border is contained in a bag of \mathcal{T} , and such that $\Lambda_{/A}$ and $\Lambda_{/B}$ are smaller than Λ . Such a partition is *good*. By induction, let $T_{/A}$ and $T_{/B}$ be p-trees of $(\Lambda_{/A}, \Pi_{/A})$ and $(\Lambda_{/B}, \Pi_{/B})$ of optimal width, and let T be the merge of $T_{/A}$ and $T_{/B}$. By Lemma 5, T is a p-tree of (Λ, Π) . By Lemma 2, its tree-width is $\max\{\text{tw}(T_{/A}), \text{tw}(T_{/B})\} = \max\{\text{tw}(\Lambda_{/A}), \text{tw}(\Lambda_{/B})\}$. Since, $\delta_\Lambda(\{A, B\})$ is contained in a bag of \mathcal{T} , Lemma 3 implies that $\text{tw}(\Lambda) = \max\{\text{tw}(\Lambda_{/A}), \text{tw}(\Lambda_{/B})\}$, and thus, $\text{tw}(T) = \text{tw}(\Lambda)$. We thus only have to find good partitions to complete the inductive step of our proof.

Three cases arise:

- Λ contains a troublesome edge e .

If e contains and separates all the other edges of Λ , then the partitioning star with one internal node is a p-tree of optimal width. Otherwise, there exists a set of edges A such that A_Π is a component of $(E_\Lambda \setminus \{e\})_\Pi$ and $\Lambda_{/A}$ is smaller than Λ . Since e , and thus $\delta_\Lambda(A)$, is contained in at least one bag of \mathcal{T} and since $\Lambda_{/(E_\Lambda \setminus A)}$ is also smaller than Λ then $\{A, E_\Lambda \setminus A\}$ is good.

- \mathcal{T} contains at least two nodes.

In any tree-decomposition of Λ with no bag being a subset of a neighbouring one, the intersection of two neighbouring bags is a separator of Λ . There thus exists a separator S which is contained in a bag of \mathcal{T} .

Let C be a component of $\Lambda \setminus S$, and let E_C be the sets of edges which are incident with vertices in C . The set E_C is Π -connected. Let $\Pi_{E_1}, \dots, \Pi_{E_p}$ be the components of $\Pi_{E \setminus E_C}$. Since $S' := \delta_\Lambda(E_C) \subseteq S$ is a separator, then there exists a component D of $\Lambda \setminus S'$ which is not C . The set of edges which are incident with D is Π -connected, and is thus a subset of, say, E_1 . Since the sets $\overline{\Pi_{E_i}}$ ($2 \leq i \leq p$) are incident with Π_{E_C} , then $\mu := \{E_1, E \setminus E_1\}$ is Π -connected. Since both $\Lambda_{/E_1}$ and $\Lambda_{/(E_\Lambda \setminus E_1)}$ contain fewer vertices than Λ and are thus smaller than Λ , and since $\delta_\Lambda(\mu) \subseteq S' \subseteq S$ is contained in a bag of \mathcal{T} , then μ is good.

- \mathcal{T} is the trivial decomposition of Λ with one node.

If Λ contains at most three edges, then all partitioning trees are p-trees, so we can suppose that Λ contains at least four edges. Since no edge of

Λ is troublesome, G^Π contains at least two vertices of degree at least 2. And since G^Π is connected, it contains at least two disjoint edges which can be extended in a non trivial bipartition of G^Π in two connected sets. This partition corresponds to a non trivial Π -connected partition $\mu := \{A, B\}$ of E_Λ . Since $\Lambda_{/A}$ and $\Lambda_{/B}$ contain fewer edges than Λ and are thus smaller than Λ , then μ is good.

□

3.5 P-trees and duality

We now turn to the second step of our proof, that is, the tree-width of a p-tree and that of its dual cannot differ too much.

Let us recall the strategy outlined in the sketch of the planar case. We prove that if T is a p-tree, then for every node v of T , the corresponding bags X_v in T and X_v^* in T^* have roughly the same size. To do so, if v is an internal node of T , then we construct a 2-cell embedding Λ' with vertex set X_v and with face set X_v^* . We then apply Euler's formula on the graph of the incidence relation between vertices and edges of Λ' to obtain our bound. It is easy to define an embedding with vertex set X_v . Indeed, if the node partition of v is $\{A, B, C\}$, then $((\Lambda_{/A})_{/B})_{/C}$ will do. The problem is that this embedding is not 2-cell, and that its face set need not be X_v^* . We thus have to be more careful when contracting to obtain such an embedding. The only problem is that we may have to consider more than 3 subsets of Σ to realise λ_v . For example, suppose that Σ is the torus and that A_Π is a cylinder. We may be forced to replace A_Π by 2 discs that fill the holes of $\Sigma \setminus A_\Pi$ and consider both discs in our realisation of λ_v . This increases the number of parts of Σ which we must consider and could cause problems when we try to bound $|X_v^*|$. But fortunately this always comes from a decrease in the genus of the surface and the required inequality remains true.

If μ is a partition of E_Λ , then we denote by $\delta_\Lambda^*(\mu)$ the set of faces of Λ which are incident with edges in at least two parts of μ .

Lemma 6 *Let Λ be a 2-cell embedding of a hypergraph on a surface Σ with at least three edges, let Π be a radial embedding of Λ , and let $\mu = \{A, B, C\}$ be a Π -connected partition of E_Λ .*

There exists a 2-cell embedding Λ' on a surface Σ' , and a partition $\mu' = \{A', B', C'\}$ of $E_{\Lambda'}$ such that

$$i. \quad V_{\Lambda'} = \delta_\Lambda(\mu), \quad F_{\Lambda'} = \delta_\Lambda^*(\mu) = \delta_{\Lambda'}^*(\mu');$$

ii. $k_{\Sigma'} \leq k_{\Sigma} + 3 - |E_{\Lambda'}|$.

Proof. Since A_{Π} , B_{Π} and C_{Π} are disjoint, we can work independently in each one of them. We thus start with A_{Π} .

First we claim that we may assume that $\Lambda \cap A_{\Pi}$ is a connected subset of A_{Π} . Since A is Π -connected, it corresponds to a connected subgraph of G^{Π} . If all the pairs $(e, f) \in A^2$ of Π -connected edges are such that $e \cup f \cap A_{\Pi}$ is connected, then we are done. So let $(e, f) \in A^2$ Π -connected edges with $e \cup f \cap A_{\Pi}$ disconnected. This can happen if e and f are incident with a vertex v in $\delta_{\Lambda}(A)$. In this case, let D be a small disc around v . We add a new vertex v_e in $e \cap D$, a new vertex v_f in $f \cap D$, and an edge $v_e v_f$ in $D \cap A_{\Pi}$. When doing so, we split a face of Λ which may belong to $\delta_{\Lambda}^*(A)$ but the new triangle face $vv_e v_f$ is only incident with edges in Π_A , and the other face belongs to $\delta_{\Lambda}^*(A)$ if and only if the original face did. The claim follows.

Let Γ be the incidence embedding of Λ . Since $\Lambda \cap A_{\Pi}$ is connected, we can contract a spanning tree of $\Gamma \cap A_{\Pi}$ so that it contains a single vertex v_A linked to $\delta_{\Lambda}(A)$ together with some loops. We then remove all the loops. Let e be such a loop. Two cases arise:

- The loop e bounds one face F whose boundary is $v_A P_1 v_A e v_A P_2$. We remove $F \cup e$ from Σ and add a disc whose boundary is $v_A P_1 v_A P_2$. If $F \cup e$ was a disc, then e would split this disc in two parts which is not the case. The genus of the new surface is thus lower than the genus of the previous one.
- The loop e bounds two faces F_1 and F_2 whose respective boundaries are $v_A P_1 v_A e$ and $v_A P_2 v_A e$. Note that $e \cup \{v_A\}$ may bound an empty disc in Σ in which case P_1 or P_2 is empty. If either F_1 or F_2 does not belong to $\delta_{\Lambda}^*(A)$, then we remove e and merge F_1 and F_2 . Otherwise, e bounds no disc. In this case, we remove e from Λ and cut Σ along $e \cup \{v_A\}$. More precisely, we remove $e \cup \{v_A\}$ from Σ and take the closure of the resulting topological space. We thus obtain a surface of lower genus and whose boundary is made of two holes $v_A^1 e^1$ and $v_A^2 e^2$. We then fill the holes with two discs. Note that in the process, v_A has been split in v_A^1 and v_A^2 but, as already mentioned the genus of the surface has dropped.

In the end, we have removed all the faces which were only incident with edges in A , and kept all the others. We have replaced A by, say, p edges and done some surgery on Σ to keep the embedding 2-cell and the surgery resulted in a decrease of genus of at least $p - 1$. The lemma follows. \square

We are now ready to prove:

Theorem 3 *Let Λ be a 2-cell embedding of a hypergraph with at least one edge on a surface Σ and let Π be a radial embedding of Λ .*

For any p -tree T of (Λ, Π) ,

$$\text{tw}(T^*) \leq \max\{\text{tw}(T) + 1 + k_\Sigma, \alpha_{\Lambda^*} - 1\}.$$

Proof. Let v be a vertex of T , let X_v be its bag in T and let X_v^* be its bag in T^* . If v is a leaf labelled by an edge e , then $X_v^* = e^*$ and $|X_v^*| - 1 \leq \max\{\text{tw}(T) + 1 + k_\Sigma, \alpha_{\Lambda^*} - 1\}$. If v is the neighbour of a leaf labelled by a troublesome edge e , then the fact that λ_v is the e -partition implies that $X_v^* \subseteq e^*$ and $|X_v^*| - 1 \leq \max\{\text{tw}(T) + 1 + k_\Sigma, \alpha_{\Lambda^*} - 1\}$.

We can thus suppose that v is an internal node of T whose node-partition $\lambda_v = \{A, B, C\}$ is Π -connected. We then have $X_v = \delta_\Lambda(\lambda_v)$, and $X_v^* = \delta_\Lambda^*(\lambda_v)$. Let Λ' be given by Lemma 6 with $\mu = \lambda_v$. Let Γ be the incidence graph of Λ' . We claim that any face of Γ is incident with at least 4 edges. This follows from the faces that Γ is bipartite, and that no face of Γ is incident with only one part A, B and C . Let F_{2k} denotes the set of faces of length $2k$. If an edge is incident with only one face F , then it counts twice in the length of F . We have $2|E_\Gamma| = 4|F_4| + 6|F_6| + \dots \geq 4|F_\Gamma|$, and thus since the face set of Γ is exactly $\delta_\Lambda^*(\lambda_v) = X_v^*$,

$$|E_\Gamma| \geq 2|X_v^*|. \quad (1)$$

Since Γ has $|X_v^*|$ faces, $|X_v| + |E_{\Lambda'}|$ vertices, and since $k_{\Sigma'} \leq k_\Sigma + 3 - |E_{\Lambda'}|$, by replacing these in Euler's formula, we obtain:

$$|X_v| + |E_{\Lambda'}| - |E_\Gamma| + |X_v^*| \geq 2 - k_\Sigma - 3 + |E_{\Lambda'}|. \quad (2)$$

Adding (1) and (2), we get

$$|X_v| + 1 + k_\Sigma \geq |X_v^*|$$

which proves that $|X_v^*| - 1 \leq \max\{\text{tw}(T) + 1 + k_\Sigma, \alpha_{\Lambda^*} - 1\}$, and thus $\text{tw}(T^*) \leq \max\{\text{tw}(T) + 1 + k_\Sigma, \alpha_{\Lambda^*} - 1\}$. \square

Theorem 1, which we restate below, is a direct corollary of Theorem 2 and Theorem 3.

Theorem 1 *For any 2-cell embedding of a hypergraph Λ on a surface Σ ,*

$$\text{tw}(\Lambda^*) \leq \max\{\text{tw}(\Lambda) + 1 + k_\Sigma, \alpha_{\Lambda^*} - 1\}.$$

4 Examples of graphs attaining the bound

In Theorem 1, the $\alpha_{\Lambda^*} - 1$ part is clearly necessary. Indeed, let Λ^* be the plane embedding of the hypergraph with one edge that contains k vertices. The dual Λ contains exactly one vertex, and thus $\text{tw}(\Lambda) = 0$ and $k - 1 = \text{tw}(\Lambda^*) = \alpha(\Lambda^*) - 1$. We now focus on embedding of graphs.

A graph G is *minimally embeddable* on a surface Σ if there exists an embedding of G in Σ and for every embedding of G on a surface Σ' , then $k_{\Sigma'} \geq k_{\Sigma}$. We also say that Γ is a minimum genus embedding on Σ . In this section, we prove the following theorem:

Theorem 4 *For any surface Σ , there exists a minimum genus embedding Γ on Σ such that $\text{tw}(\Gamma) = \text{tw}(\Gamma^*) + 1 + k_{\Sigma}$.*

To do so, we need some more definitions. A *bramble* of G is a family \mathcal{B} of subsets of vertices of G such that for every element $X \in \mathcal{B}$, $G[X]$ is connected, and for all elements $X, Y \in \mathcal{B}$, $G[X \cup Y]$ is connected (we say that X and Y *touch*). The *order* of a bramble \mathcal{B} is the minimum size of a set $X \subseteq V_G$ which intersects all the elements of \mathcal{B} . We use brambles to compute tree-width with the following theorem:

Theorem 5 ([ST93]) *The maximum order of a bramble of G is $\text{tw}(G) + 1$.*

Let Γ be an embedding of a graph on a surface Σ which is not the sphere, and let $\theta > 0$ be an integer. We say that Γ is *θ -representative* if for every $\mu \subseteq \Sigma$ homeomorphic to a circle, if γ is not the boundary of some closed disc in Σ (we call such μ a *non-contractible noose*) then $|\gamma \cap \Gamma| \geq \theta$. We use representativity to certify that graphs are minimally embedded with the following theorem which is an easy corollary of a theorem in [ST96].

Theorem 6 *Let Γ be a θ -representative embedding of a graph on a surface Σ which is not the sphere. If $\theta \geq 100k_{\Sigma}$ then Γ is a minimum genus embedding.*

4.1 Todinca graphs and their tree-widths

Let $p \geq 1$ be an integer. Let A, B and C be three $2p \times 2p$ grids. Let $a_1, \dots, a_p, a'_p, \dots, a'_1$ be the top row of A . Let $b_1, \dots, b_p, b'_p, \dots, b'_1$ be the top row of B . Let $c_1, \dots, c_p, c'_p, \dots, c'_1$ be the top row of C . A *Todinca graph*⁴ of order p is any graph G obtained by bijectively linking the edges between a_1, \dots, a_p and b'_1, \dots, b'_p , between b_1, \dots, b_p and c'_1, \dots, c'_p and between c_1, \dots, c_p and a'_1, \dots, a'_p . As an example, in Figure 5, the left graph is obtained

⁴The name come from Ioan Todinca who first showed us the plane graph of Figure 5 as an example of graph whose tree-width + 1 equals 3/2 its branch-width.




Figure 5: Two Todinca graphs.

by adding the edges $a_i b'_i$, $b_i c'_i$ and $c_i a'_i$, while in the right graph, we link b_i to c'_{p+1-i} .

Lemma 7 *The tree-width of every Todinca graph G of order p is at least $3p - 1$.*

Proof. By Theorem 5, we only have to produce a bramble of order $3p$ to prove our lemma. To do so, we give some definitions.

The edges added to the grids A , B and C link their columns and thus define *columns* of G . Depending on which columns were linked, we obtain *AB-columns*, *BC-columns* and *CA-columns*. We also call the respective rows of A , B and C , the *A-rows*, the *B-rows* and the *C-rows* of G . Together, they are the *rows* of G (see Figure 6). The union of an *A-row* and an *AB-column* is an *A-cross*. In the same spirit, we define *B-crosses* and *C-crosses*, and we claim that the set \mathcal{C} of all these *crosses* is a bramble of G of order $3p$.

They clearly are connected, so let us prove that any two crosses X and Y touch. By symmetry, we can suppose that X is an *A-cross*. Since X contains an *A-row*, and since *A-rows* intersect all *AB-* and all *CA-columns*, if Y is an *A-* or a *C-cross*, then X and Y touch. If Y is a *B-cross*, then it contains a *B-row* which intersects all *AB-columns*. The crosses X and Y therefore touch, which finishes our proof that \mathcal{C} is a bramble.

Since there are exactly $3p$ columns, the order of \mathcal{C} is at most $3p$. Let us prove that no set S of size $3p - 1$ can intersect all the crosses. Obviously, the set $S \subseteq V_G$ misses at least one column. We split this proof in two cases

- There exists an *AB-column* X , a *BC-column* Y and a *CA-column* Z avoiding S . Since G contains $6p$ rows, there exists a row L avoiding S .




Figure 6: The rows and the columns of a Todinca graph.

Depending on what kind of row L is, one of $L \cup X$, $L \cup Y$ and $L \cup Z$ is a cross of G avoiding S .

- There exists, by symmetry, an AB -column C avoiding S , and no BC -column avoiding S . Since there are at least p vertices of S on the BC -columns, there can only be $2p - 1$ vertices from S on the grid A . There thus exists an A -row L avoiding S . The cross $L \cup C$ avoids S .

□

We also need the following folklore lemma (see Figure 7).

Lemma 8 *An $n \times m$ grid G has tree-width at most $\min(n, m)$, and G admits a path-decomposition attaining this bound with a leaf containing one of its shortest sides.*




Figure 7: A tree-decomposition of the 4×5 grid of width 4.

As a matter of fact, the tree-width of an $n \times m$ grid G is exactly $\min(n, m)$ (see [BGK08]).

Lemma 9 *The tree-width of every Todinca graph G of order p is at most $3p - 1$.*

Proof. To prove the lemma, we give a tree-decomposition of G of width $3p - 1$. Let A , B and C be the three $2p \times 2p$ grids of G , and let \mathcal{T}_A , \mathcal{T}_B and \mathcal{T}_C be path decompositions of A , B and C given by Lemma 8. Let v_A , v_B and v_C be the vertices of \mathcal{T}_A , \mathcal{T}_B and \mathcal{T}_C whose bags respectively contain the top rows of A , B and C . Let u , u_A , u_B and u_C be vertices whose bags respectively are $\{a_1, \dots, a_p, b_1, \dots, b_p, c_1, \dots, c_p\}$, $\{a_1, \dots, a_p, a'_1, \dots, a'_p, b_1, \dots, b_p\}$, $\{b_1, \dots, b_p, b'_1, \dots, b'_p, c_1, \dots, c_p\}$ and $\{c_1, \dots, c_p, c'_1, \dots, c'_p, a_1, \dots, a_p\}$. The labelled tree T obtained by linking u to u_A , u_B and u_C , and adding the edges $u_A v_A$, $u_B v_B$ and $u_C v_C$ is a tree-decomposition of G of width $3p - 1$. \square

Lemmata 7 and 9 clearly imply that

Lemma 10 *The tree-width of every Todinca graph of order p is $3p - 1$.*

4.2 Some minimally embeddable Todinca graphs

We now define special Todinca graphs. We first define three gadgets (see Figure 8) that we use to link the grids of our graphs. Let P and Q be two

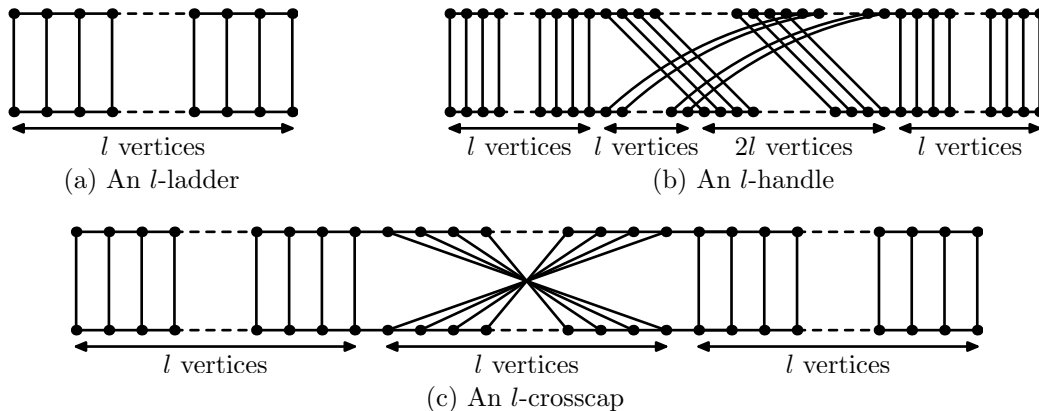


Figure 8: linking gadgets

paths of length k , and let p_1, \dots, p_k and q_1, \dots, q_k be their respective vertices. If we link each p_i with q_i ($1 \leq i \leq k$), we obtain a k -ladder. If $k = 3l$ and we link p_i with q_i ($1 \leq i \leq l$ and $2l + 1 \leq i \leq 3l$) and p_{l+i} with q_{2l+1-i} ($1 \leq i \leq l$), then we obtain an l -crosscap. If $k = 5l$ and we link p_i with q_i ($1 \leq i \leq l$ and $4l + 1 \leq i \leq 5l$), p_{l+i} to q_{2l+i} ($1 \leq i \leq 2l$) and p_{3l+i} to q_{l+i} ($1 \leq i \leq l$), then we obtain an l -handle.

Let $k > 0$ and $p > 1$ be integers. The graph $G_{k,p}$ is the Todinca graph of order $5kp$ such that the grid A is linked to the grids B and C by a $5kp$ -ladder, and the grids B and C are linked k p -handles. The graph $\tilde{G}_{k,p}$ is the Todinca

graph of order $3kp$ defined in the same way except that we use p -crosscaps instead of p -handles.

We now define an embedding $\Gamma_{k,p}$ of $G_{k,p}$ in Σ_k as follows. We start by embedding the three grids and the $5kp$ -ladders on the sphere. We also embed the all the edges of the handles except the edges $p_{3l+i}q_{l+i}$ ($1 \leq i \leq l$). Then for each p -handle, we add a handle to the surface to embed the remaining p edges. We also define an embedding $\tilde{\Gamma}_{k,p}$ of $\tilde{G}_{k,p}$ in $\tilde{\Sigma}_k$ in a similar way except that for each p -crosscap, we add a crosscap to the surface so that the edges $p_{l+i}q_{2l+1-i}$ ($1 \leq i \leq l$) do not cross.

We now want to prove that for p large enough, these embeddings are minimal genus embeddings. To do so, we want to apply Theorem 6. We thus have to prove that for p large enough, these embeddings have large representativity. More precisely,

Lemma 11 *The embeddings $\Gamma_{k,p}$ and $\tilde{\Gamma}_{k,p}$ are p -representative.*

Proof. The result for $\Gamma_{k,p}$ and $\tilde{\Gamma}_{k,p}$ are respectively direct consequences of Theorems 3.5 and 3.3 of [RS96], but very few people seem to have read the graph minors paper so we give a direct proof which is, in spirit, very similar to the proofs by Robertson and Seymour.

We start with $\Gamma_{k,p}$. As already said, to embed $G_{k,p}$ on Σ_k , we start by embedding a subgraph of $G_{k,p}$ on the sphere, then for each handle gadget P_i ($1 \leq i \leq k$), we add a handle to embed the edges $p_{3l+i}q_{l+i}$ ($1 \leq i \leq l$). To do so, we remove two faces f_i and g_i of our partial embedding and we “sew” a cylinder H_i to the border of the holes. It is easy to find p vertex disjoint concentric cycle c_1^i, \dots, c_p^i enclosing f_i and p vertex disjoint concentric cycles d_1^i, \dots, d_p^i enclosing g_i such that all the circles are pairwise disjoint (see left part of Figure 9). The component of $\Sigma_k \setminus (\bigcup_{i=1}^k (c_p^i \cup d_p^i))$ which is incident




Figure 9: Enclosing concentric cycles around a handle or a crosscap.

with all the cycles c_p^i and d_p^i is the *outside region*. Let μ be a non-contractile noose on Σ_k . We split our proof in the following three sub-cases.

- Suppose that μ intersects both the outside region and the handle H_i . The curve μ has to cross all the cycles c_1^i, \dots, c_p^i or all the cycles d_1^i, \dots, d_p^i to reach the outside region, which implies that $|\mu \cap \Gamma_{k,p}| \geq p$.
- Suppose that μ intersects a handle H_i . Then μ is a subset of the component H'_i of $\Sigma_k \setminus (c_p^i \cup d_p^i)$ which contains H_i . It is easy to find p vertex distinct paths Q_1, \dots, Q_p in $\Gamma_{k,p}$ which link vertices in c_p^i and vertices in d_p^i and whose interior is in H'_i (see left part of Figure 10).




Figure 10: Paths across a handle or a crosscap.

We claim that μ intersects all the paths Q_i . Indeed otherwise μ is a subset of $H'_i \setminus Q_j$ for some $1 \leq j \leq p$. But $H'_i \setminus Q_j$ is an open disc which implies that μ is contractible, a contradiction. This thus implies that $|\mu \cap \Gamma_{k,p}| \geq p$.

- Suppose that μ intersects the outside region. Then μ is a subset of $\Sigma' = \Sigma_k \setminus (\bigcup_{i=1}^k H_i)$. Since Σ' is a sphere with holes bounded by the cycles c_1^i and d_1^i , then μ separates c_1^1 and c_1^1 or d_1^1 in Σ' . In both cases, it is easy to find p vertex disjoint paths in $\Gamma_{k,p}$ between c_1^1 and c_1^1 or d_1^1 whose interior are in Σ' (see upper part of Figure 11). Since μ has to cut these paths, $|\mu \cap \Gamma_{k,p}| \geq p$.

In all cases, $|\mu \cap \Gamma_{k,p}| \geq p$ which finishes our proof that the embedding of $\Gamma_{k,p}$ is p -representative.

The proof that $\tilde{\Gamma}_{k,p}$ is p -representative is very similar. We first enclose each crosscap by p disjoint cycles, and then we define an outside region. We then prove that any noose μ which intersects both the outside and a crosscap meets the embedding at least p times. If μ intersects a crosscap, it is enclosed by the outer cycle c enclosing the crosscap. We can easily find p disjoint paths each linking two points of c as in the right part of Figure 10, and μ has to meet all those paths otherwise it is contractible. If μ intersects the outside region, then it must separate some crosscaps and since it is easy to find p disjoint paths linking the inner cycles enclosing these crosscaps, μ intersects the embedding at least p times. \square




Figure 11: Paths linking handles or crosscaps in the outside region.

As a consequence of Lemma 11 and Theorem 6, we have

Lemma 12 $\Gamma_{k,100k}$ and $\tilde{\Gamma}_{k,100k}$ are minimum genus embeddings respectively in Σ_k and $\tilde{\Sigma}_k$.

4.3 Dual of some Toinca graphs

Lemma 13 The tree-width of $\Gamma_{k,p}^*$ is at most $15kp - 2 - k$.

Proof. The embedding $\Gamma_{k,p}^*$ is made of three $(10kp - 1) \times (10kp - 1)$ grids A , B and C , two paths of $5kp - 1$ vertices P_{AB} and P_{AC} , a graph G_{BC} which corresponds to the dual of the gadgets and two vertices v_{in} and v_{out} . The path P_{AB} is adjacent to the grids A and B , the path P_{AC} is adjacent to the grids A and C , and the graph G_{BC} is adjacent to the grids B and C . The vertex v_{out} is adjacent to the vertices of the bottom rows of A , B and C , to the vertices of the side columns of A , B and C and to the “outer vertex” of P_{AB} , P_{AC} and G_{BC} . The vertex v_{in} is adjacent to the middle vertex of the top rows of A , B , C and to the “inner vertex” of P_{AB} , P_{AC} and G_{BC} (see Figure 12).

$\Gamma_{k,p}$ is an embedding of a Toinca graph of order $l = 5kp$ in a surface of Euler genus $2k$. It thus has $12l^2$ vertices and $24l^2 - 9l$ edges. By Euler’s formula, it thus has $12l^2 - 9l + 2 - 2k$ faces. This number is also the number of vertices of $\Gamma_{k,p}^*$. There are $3(2l - 1)^2$ vertices in the grids, $2(l - 1)$ vertices on P_{AB} and P_{AC} and the two vertices v_{in} and v_{out} . This leaves $l - 2k - 1$ gadget vertices on G_{BC} . Since the tree-width of an embedding drops by at most one

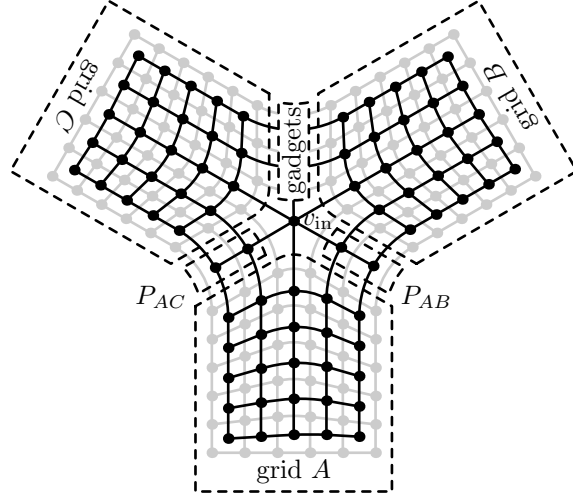


Figure 12: The embedding $\Gamma_{k,p}^* - v_{\text{out}}$.

when removing a single vertex, we only have to prove that the tree-width of $\Gamma_{k,p}^* - v_{\text{out}}$ is $15kp - 3 - 2k$ so let us remove the vertex v_{out} from $\Gamma_{k,p}^*$.

The grid A together with its neighbourhood is a $2l \times (2l - 1)$ grid. We can thus choose a path decomposition \mathcal{T}_A of width $2l - 1$ of this grid and in which a vertex v_A contains the neighbourhood of A . Because of the gadgets, the links between the grid B with its neighbourhood is more complex. Let v_1, \dots, v_{2l-1} be the vertices of the top row of B which link B to the remaining of $\Gamma_{k,p}^*$. The vertices v_1, \dots, v_l are clearly linked to $P_{AB} \cup \{v_{\text{in}}\}$ in an l -ladder. Since the gadget sequence begins with a p -ladder, there is a set S of $p - 1$ gadget vertices such that v_1, \dots, v_{l+p-1} is linked to $P_{AB} \cup \{v_{\text{in}}\} \cup S$ in a $l + p - 1$ ladder. The remaining gadget vertices are linked to the vertices v_{l+p}, \dots, v_{2l-1} . Using the idea behind the path decomposition of Figure 7, we define a path decomposition \mathcal{T}_B of width $3l - p - 2k - 2$ of B and its neighbourhood in which a vertex v_B contains P_{AB}, v_{in} , the gadget vertices and v_{l+p}, \dots, v_{2l-1} . We can similarly define a tree-decomposition \mathcal{T}_C of C and its neighbourhood in which a vertex v_C contains the neighbourhood of C . Let u be a vertex whose bag contains $P_{AB}, P_{AC}, v_{\text{in}}$ and the gadget vertices, and let us add the edges uv_A, uv_B and uv_C . This defines a tree-decomposition of $\Gamma_{k,p}^* - v_{\text{out}}$. The bag of u is its biggest one and it contains $3l - 2 - 2k$ vertices which proves that the tree-width of $\Gamma_{k,p}^*$ is at most $3l - 2 - 2k = 15kp - 2 - 2k$. \square

Using a similar proof, we also obtain:

Lemma 14 *The tree-width of $\tilde{\Gamma}_{k,p}^*$ is at most $9kp - 2 - k$.*

We are now ready to prove Theorem 4.

Theorem 4 *For any surface Σ , there exists a minimum genus embedding Γ on Σ such that $\text{tw}(\Gamma) = \text{tw}(\Gamma^*) + 1 + k_\Sigma$.*

Proof. In [RS84], Robertson and Seymour already gave examples of planar embeddings matching our bound. So let us consider higher genus surfaces.

$\Gamma_{k,100k}$ is an embedding of a Todinca graph of order $500k^2$. By Lemma 10, its tree-width is $1500k^2 - 1$. By Lemma 13, the tree-width of $\Gamma_{k,100k}^*$ is at most $1500k^2 - 2 - 2k$. Since $\Gamma_{k,100k}$ is an embedding in Σ_k , Theorem 1 implies that $\text{tw}(\Gamma_{k,100k}^*) = 1500k^2 - 2 - 2k$ and $\text{tw}(\Gamma_{k,100k}) - \text{tw}(\Gamma_{k,100k}^*) = 1 + 2k$. Since, by Lemma 12, $\Gamma_{k,100k}$ is a minimum genus embedding in Σ_k , $\Gamma_{k,100k}$ and $\Gamma_{k,100k}^*$ indeed are examples of embeddings matching the bound of Theorem 1 for the surface Σ_k .

Similarly, $\tilde{\Gamma}_{k,100k}$ and $\tilde{\Gamma}_{k,100k}^*$ are examples of embeddings matching the bound of Theorem 1 for the surface $\tilde{\Sigma}_k$. \square

5 Conclusion and open questions

In this paper, we show that tree-width is a robust parameter considering surface duality. Indeed, our main proof says more than just “the difference between the tree-width of Λ and that of Λ^* is small”. Our proof says that there always exists a decomposition which is optimal for Λ and very good for Λ^* . This leads to a natural question: For any embedding Λ , does there always exist a p-tree T such that $\text{tw}(T) = \text{tw}(\Lambda)$ and $\text{tw}(T^*) = \text{tw}(\Lambda^*)$? To our knowledge, the question is open, even for plane embeddings of graphs.

References

- [AP89] Stephan Arnborg and Andrzej Proskurowski. Linear time algorithms for NP-hard problems restricted to partial k -trees. *Discrete Applied Mathematics*, 23:11–24, 1989.
- [BGK08] Hans Leo Bodlaender, Alexander Grigoriev, and Arie Marinus Catharinus Antonius Koster. Treewidth lower bounds with brambles. *Algorithmica*, 51:81–98, 2008.
- [BMT03] Vincent Bouchitté, Frédéric Mazoit, and Ioan Todinca. Chordal embeddings of planar graphs. *Discrete Mathematics*, 273:85–102, 2003.

- [Cou90] Bruno Courcelle. The monadic second-order logic of graphs I: recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, 1990.
- [Hal76] Rudolf Halin. S -functions for graphs. *Journal of Geometry*, 8:171–186, 1976.
- [Lap96] Denis Lapoire. Treewidth and duality for planar hypergraphs. Manuscript http://www.labri.fr/perso/lapoire/papers/dual_planar_treewidth.ps, 1996.
- [Maz04] Frédéric Mazoit. *Décomposition algorithmique des graphes*. PhD thesis, École normale supérieure de Lyon, 2004.
- [Maz09] Frédéric Mazoit. Tree-width of graphs and surface duality. *Electronic Notes in Discrete Mathematics*, 32:93–97, 2009.
- [RS84] Neil Robertson and Paul D. Seymour. Graph Minors. III. Planar Tree-Width. *Journal of Combinatorial Theory Series B*, 36(1):49–64, 1984.
- [RS96] Neil Robertson and Paul D. Seymour. Graph Minors. XV. Giant Steps. *Journal of Combinatorial Theory Series B*, 68(1):112–148, 1996.
- [ST93] Paul D. Seymour and Robin Thomas. Graph Searching and a Min-Max Theorem for Tree-Width. *Journal of Combinatorial Theory Series B*, 58(1):22–33, 1993.
- [ST96] Paul D. Seymour and Robin Thomas. Uniqueness of Highly Representative Surface Embeddings. *Journal of Graph Theory*, 23(4):337–349, 1996.