International Symposium on Theoretical Aspects of Computer Science, STACS 2012, February 29 - March 3, 2012.

Contraction checking in graphs on surfaces

Marcin Kamiński[1⋆] and Dimitrios M. Thilikos[2⋆⋆]

[1] Département d'Informatique
Université Libre de Bruxelles
Marcin.Kaminski@ulb.ac.be

[2] Department of Mathematics
National and Kapodistrian University of Athens
sedthilk@math.uoa.gr

Abstract. The Contraction Checking problem is to decide whether H can be obtained from G by a sequence of edge contractions. Contraction Checking remains NP-complete, even when the size of H is fixed. We prove that for every surface Σ, there exists a cubic algorithm for Contraction Checking, when the input graph is Σ-embeddable.

⋆ Chargé de Recherches du FRS-FNRS.

⋆⋆ Supported by the project "Support of Research" of the National and Kapodistrian University of Athens (project code: 10812).

1 Introduction

We consider simple finite graphs and use standard graph-theoretical terminology. For notions not define here, we refer the reader to Diestel [6] and to Mohar and Thomassen [19].

Contractions and topological minors. To *contract an edge* is to identify its two endpoints and remove the loop and multiple edges that have possibly been created. A graph H is a *contraction* of a graph G ($H <_c G$) if H can be obtained from G by a sequence of edge contractions. Deciding whether the input graph can be contracted to a fixed pattern is NP-complete, even for small pattern graphs – the smallest is an induced path on four vertices [3].

To *dissolve a vertex* of degree 2 is to contract one of the edges incident with it. A graph H is a *topological minor* of a graph G if H can be obtained from G by a sequence of vertex/edge deletions and vertex dissolutions. Recently, Grohe et al. proved that for every fixed graph H there exists an $\mathcal{O}(|V(G)|^3)$ time algorithm deciding whether H is a topological minor of G [14]. This is an FPT algorithm for this problem when parameterized by the size of H, that is, an algorithm with running time $g(|H|) \cdot |G|^{\mathcal{O}(1)}$. (For more information on parametrized complexity theory, see any of the books: Downey and Fellow [7], Flum and Grohe [10], or Niedermeier [22].)

Surface containment relations. Surface versions of contractions and topological minors can be defined for surface-embedded graphs. Formal definitions are presented in Section 3. For the purpose of this introduction, we only note that surface contractions and surface topological minors are surface-embedded versions of contractions and topological minors, respectively, that respect the embedding.

For every surface Σ and every pattern graph H, there exists a polynomial-time algorithm deciding whether a Σ-embedded graph can be contracted to H [15]. The algorithm is based on a combinatorial lemma that allows to reduce the problem of testing for contraction in a surface-embedded graph to a constant number of tests for surface topological minors in its dual. The procedure is polynomial for every fixed graph H; however, the degree of the polynomial depends on the size of H. Is it possible to design an FPT algorithm for this problem when parameterized by the size of H?

The main obstacle is testing for surface topological minors. If there existed an FPT algorithm for deciding if a surface-embedded input graph contains a pattern graph H as a surface topological minor, then the machinery of [15] would imply a FPT algorithm for contraction checking. Surface topological minors are different from topological minors as they are defined for surface-embedded graphs and respect the embedding. While it is possible to reduce topological minor testing to surface topological minor testing, the latter is not known to be FPT-reducible to the former.

In this paper we overcome these difficulties and show that testing whether a surface-embedded graph is contractible to a given pattern is FPT, when parameterized by the size of the pattern.

The irrelevant vertex technique. A core technique from Graph Minors by Robertson and Seymour that has been especially prolific in algorithmic research is the following win/win approach. If the treewitdh of the input graph is small (less than a certain constant c), apply dynamic programming and solve the problem in FPT time with respect to c; otherwise, exploit the existence of a subdivision of a large wall in the input graph (its size depends on c). In the latter case, one can usually find an *irrelevant vertex* – a

Relation	planar graph	graphs on surfaces	all graphs
(induced) subgraph	FPT [8]		W[1]-hard
minor	FPT, [24]		
induced minor	FPT [9]	OPEN	XP-hard [9]
contraction	FPT [this paper]		XP-hard [3]
topological minor	FPT [14]		
weak/strong immersion	FPT [14]		

Table 1. Overview of parameterized complexity status of containment relations in graphs.

vertex that can be safely removed from the graph without changing the solution. Then, the algorithm is recursively applied to the new graph so that, eventually, the treewidth of the graph drops below c to make the dynamic programming approach applicable.

Our approach. We follow this general scheme, however, we additionally prove that one can assume that the subgraph containing a large subdivided wall is of bounded treewidth. More precisely, for every positive integer h and a surface Σ, there exist constants t and T such that in every Σ-embedded graph of treewidth at least t there exists a disk in Σ such that the graph induced by the vertices inside the disk is of treewidth at most T and contains a subdivision of a wall of height h. This assumption comes in handy in our proof. We also believe that this lemma is of independent interest and can also be applied to other problems.

Having found a subgraph of bounded treewidth containing a large subdivided wall, we consider a collection of nested cycles from the wall.For each cycle from the collection, we check what sub-patterns of the guest graph can be seen as surface topological minors of its interior with a "certain attachment" to the boundary of the cycle. This attachment determines the possible ways such a pattern should be extended outside the cycle towards matching the structure of the host graph. This is encoded as a *characteristic* function of each cycle. A key property is that the characteristic function is *monotone* – whatever can be attached to a cycle, can also be attached to subsequent cycles in the collection.

The main idea is to determine a collection of consecutive cycles with the same characteristic function, which is now fesible since this computation takes place in a graph of bounded treewidth. If this collection is "sufficiently large" then the monotonicity property implies that every sub-pattern of the guest graph can be also located away from some "safe" cycle and this is is proved by making use of the Unique Linkage Theorem of Robertson and Seymour in [23,25]. Then the safe cycle contains an irrelevant vertex that is removed and the procedure recurses until the host graph has bounded treewidth.

We show that contractions are FPT for graphs embedded on surfaces. Table 1 summarizes the current state of research on parameterized complexity of containment relations.

2 Previous work on contractions

The problem of checking whether a graph is a contraction of another has already attracted some attention.

Perhaps the first systematic study of contractions was undertaken by Brouwer and Veldman [3]. According to the results of [3], checking if a graph is contractible to the induced cycle on four vertices or the induced path on four vertices is NP-complete. More generally, it is NP-complete for every bipartite graph with at least one connected component that is not a star. Looking at contractions to fixed pattern graphs is justified by the following result proved by Matoušek and Thomas [17].

Proposition 1 ([17]). *The problem of deciding, given two input graphs G and H, whether G is contractible to H is* NP*-complete even if we impose one of the following restrictions on G and H: (i) H and G are trees of bounded diameter, or (ii) H and G are trees all whose vertices but one have degree at most* 5*. Moreover, for every fixed k, the problem of deciding, given two input graphs G and H, whether G is contractible to H is* NP*-complete even if we restrict G to partial k-trees and H to k-connected graphs.*

3 Definitions

Surfaces. A *surface* Σ is a compact 2-manifold without boundary (we always consider connected surfaces). Whenever we refer to a Σ-*embedded graph* G we consider G accompanied by some embedding of it in Σ without crossings. To simplify notation, we do not distinguish between a vertex of G and the point of Σ used in the drawing to represent the vertex or between an edge and the line representing it. Given an edge e, we denote by $\overline{e}$ the set of its endpoints (clearly, $1 \leq |\overline{e}| \leq 2$). We also consider a graph G embedded in Σ as the union of the points corresponding to its vertices and edges. That way, a subgraph H of G can be seen as a graph H, where $H \subseteq G$. We refer to the book of Mohar and Thomassen [20] for more details on graph embeddings. The *Euler genus* of a graph G is the minimum integer γ such that G can be embedded on a surface of the Euler genus γ.

Given a Σ-embedded graph G, we denote by $F(G)$ the set of its faces, i.e. the set of connected components of the set $\Sigma \setminus G$. We say that a face in $F(G)$ is *trivial* if it is incident with at most two edges. An edge is *trivial* if it is incident with a trivial face. A loop of G is an edge with one endpoint. We say that a loop e is *singular* if it is either non-contractible or it is contractible and both connected components of $\Sigma \setminus e$ contain vertices of G.

The *surface contraction of an edge* e in a Σ-embedded graph G is the graph $G' = G\backslash_{\Sigma} e$ defined as follows. In case e is non-singular, G' is the graph obtained if we identify the closure of all points of e to a single vertex. In case e is singular the G' is the graph obtained from G after removing all points of e. Notice that surface contractions are defined in a way that the surface integrity is maintained.

Let H and G be two Σ-embedded graphs. We say that H is a *surface contraction* of G, denoted by $H \leq_{\mathrm{c}}^{\Sigma} G$, if H can be obtained from G by a (possibly empty) sequence of operations that may be either surface contractions of edges or removals of trivial edges. Finally, we say that H is a *surface minor* of G, if H is a surface contraction of some subgraph of G.

Isomorphism. Let A_1 and A_2 be graphs and let $\psi\colon V(A_1) \to V(A_2)$ be a bijection. We say that A_1 and A_2 are ψ-*isomorphic* if for each pair $x, y \in V(A_1)$ it holds that $\{x, y\} \in E(A_1)$ if and only if $\{\psi(x), \phi(y)\} \in E(A_2)$. The *edge extension* of ψ, denoted by $\psi^e : V(A_1) \cup E(A_1) \to V(A_2) \cup E(A_2)$ extends ψ so to incorporate the correspondence between the edges of A_1 and the edges of A_2 implied by ψ.

Topological isomorphism. Let A_i be Σ_i-embedded graphs $i \in \{1, 2\}$. Suppose also that Σ_1 is homeomorphic to Σ_2. Let $\psi : V(A_1) \to V(A_2)$ be a bijection from $V(A_1)$ to

$V(A_2)$. We say that A_1 is ψ-topologically isomorphic to A_2 is there is a homeomorphism $\phi : \Sigma_1 \to \Sigma_2$ such that

- ψ is an isomorphism from A_1 to A_2 and
- ψ^e is induced by the restriction of ϕ in $V(A_1)$.

Notice that the bijection ψ above is an isomorphism between A_1 and A_2.

Surface topological minor. Let Σ be a surface and $\mathcal{G}$ be a Σ-embedded graph. Given a set $\mathcal{P}$ of internally disjoint extended paths of G, we define $G_{\mathcal{P}}$ as the Σ-embedded graph created if we first remove from G each edge not in a path in $\mathcal{P}$ and then replace each extended path (P, A) in $\mathcal{P}(G)$ by the extended path $((\overline{e}, \{e\}), A)$ where e is a new edge and $\overline{e} = A$.

Let Σ be a surface and (G, S_G) and (H, S_H) be two rooted Σ-embedded graphs. Let also σ be a bijection from S_G to S_H. We say that (H, S_H) is a *surface σ-rooted topological minor* of (G, S_G), and we denote it by $(H, S_H) \leq_{\sigma}^{\Sigma} (G, S_G)$ if there is a collection $\mathcal{P}$ of internally disjoint extended paths in G such that $G_{\mathcal{P}}$ is ψ-topologically isomorphic to H for some bijection $\psi : V(G_{\mathcal{P}}) \to V(H)$ where $\sigma \subseteq \psi$. When $S_G = S_H = \emptyset$, we say that H is a *surface topological minor* for G and denote it by $H \leq_{\text{stm}}^{\Sigma} G$.

The main technical result of [15], [16] is an equivalence between surface contractions in a surface-embedded graph and surface topological minors in its dual. A multigraph is called *thin* if it has no two parallel edges bounding a 2-face. (In particular, simple graphs are thin.) For a surface Σ and a simple Σ-embedded graph H, let $C_{\Sigma}(H)$ be a maximal set of thin Σ-embedded multigraphs that have the same adjacencies between their vertices as H (that is, forgetting multiple edges) such that they are all combinatorially different. The set $C_{\Sigma}(H)$ is finite ([15], [16]).

Proposition 2 ([15], [16]). *Let G and H be graphs. Suppose also that G is embedded in a surface Σ and let G^* be its dual. Then $H \leq_{\text{c}}^{\Sigma} G$ if and only if there exists a graph $\hat{H} \in C_{\Sigma}(H)$ such that $\hat{H}^* \leq_{\text{stm}}^{\Sigma} G^*$.*

4 Description of the algorithm

Let G and H be the host and the guest graph respectively. We denote by n the number of vertices in G. Also, in order to maintain only one parameter during the description of the algorithm, we assume that $h = |E(H)| + |V(H)| + \mathbf{eg}(G)$, where $\mathbf{eg}(G)$ is the Euler genus of G. For simplicity, we will use the notation $O_h(n^{\alpha})$ instead of $f(h) \cdot n^{\alpha}$ where f is some computable function of h.

General framework. Following the idea of the irrelevant vertex technique, introduced by Robertson and Seymour in [24], our first step is to check whether the treewidth of G is at most $f_0(H) + h + 1$ where $f_0 : \mathbb{N} \to \mathbb{N}$ is a suitable function of H. This can be done in $O_h(n)$ steps because of the results in [2]. If $\mathbf{tw}(G) < f_0(h) + h + 1$, then the problem can be solved by the dynamic programming algorithm of [1] in $O_h(n)$ steps (this also follows from Courcelle's theorem [4] and the fact that contraction checking is expressible in Monadic Second Order Logic). So we may assume that $\mathbf{tw}(G) \geq f_0(h) + h + 1$. Also using the algorithm in [21] we may consider that G is optimally 2-cell embedded in some surface Σ of Euler genus $\mathbf{eg}(G)$. Let G^* be the dual embedding of G in Σ. From [18], the treewidth of a Σ 2-cell embedded graph and the treewidth of its dual cannot differ more than $\mathbf{eg}(\Sigma)+1$. Therefore $\mathbf{tw}(G^*) \geq f_0(h)$. From Proposition 2, H is a contraction of G if and only if for some Σ-embedded graph in $\hat{H} \in C_\Sigma(H)$ it holds that $\hat{H}^* \leq_{\text{stm}}^{\Sigma} G^*$. Recall that the size of each graph in $C_\Sigma(H)$ depends only on H and $\mathbf{eg}(G)$ and therefore is bounded by $f_1(h)$ for some function f_1.

Our goal is to give an $O_h(n^2)$ step procedure with the following specifications:

Procedure Irrelevant Edge Detection(G, Σ)
Input: a graph G' of treewidth at least $f_0(h)$ that is 2-cell embedded in a surface Σ of Euler genus $\leq h$.
Output: an edge $e' \in E(G')$ such that $G' \setminus e$ remains 2-cell embedded in Σ and for every Σ-embedded graph H' of size at most $f_1(h)$, it holds that

$$H' \leq_{\text{stm}}^{\Sigma} G' \Leftrightarrow H' \leq_{\text{stm}}^{\Sigma} G' \setminus e'.$$

Actually, function f_0 should be chosen to be "sufficiently big" so it is possible to find an irrelevant edge.

Let e^* be the output of Irrelevant Edge Detection(G^*, Σ). Using the proof of Proposition 2, we may find an edge $e^* \in E(G^*)$ such that if e^* it is the dual edge of $e \in E(G)$, then H is a contraction of G if and only if H is a contraction of G/e. That way we reduce, in $O_h(n^2)$ steps, the problem of checking whether $H \leq_c G$ to the problem whether $H \leq_c G_{\text{new}} = G/e$. Clearly, we may again check whether $\mathbf{tw}(G_{\text{new}}) < f_0(h) + h + 1$ and either solve the problem by dynamic programming or again apply the Irrelevant Edge Detection procedure on G_{new}. Since the new graph is always smaller than the previous, applying the same steps, the algorithm will stop and produce a correct solution. As this will occur in less than n repetitions, the whole algorithm will take $O_h(n^3)$ steps, as claimed.

Given the above framework, what remains is to describe how the Irrelevant Edge Detection procedure works.

Big walls of small treewidth. It follows from the results in [5,12,13] that every Σ-embeddable graph of big enough treewidth contains as a subgraph a subdivision of a wall of given height and width (where height and width are defined in the obvious way). Also, by the same results, we can assume that this subdivision is "flat in the surface" in the sense that its perimeter is a contractible cycle of the embedding (i.e. handles are outside the wall). An example of such a subdivided wall is depicted in Figure 1 (for simplicity, we do not depict the subdivision vertices). We need the following Lemma:

Lemma 1. *There are functions t_1 and t_2 such that, for every κ, every graph G that is embedded in a surface Σ of Euler genus g and has treewidth at least $t_1(\kappa, g)$, contains a subgraph R such that*

- *R is the subdivision of a wall of height and width equal to k,*
- *R is drawn inside a closed disk Δ bounded by its perimeter, and*
- *$\Delta \cap G$, i.e. the part of the graph that lies inside the perimeter of R, has treewidth upper bounded by $t_2(\kappa, g)$.*

Also, such a graph R can be computed in $O_h(n^2)$ steps.

Proof. The following claim can easily be derived by Lemma 4 in [12].

Claim. Let G be a graph embedded in a surface Σ of Euler genus g and let i be a positive integer. If $\mathbf{tw}(G) \geq 48i(g+1)$, then G contains a subdivided wall R of height i and width i as a subgraph and R is drawn inside a closed disk Δ of Σ bounded by the perimeter of G'.

Let $t_1(\kappa, g) = 48\kappa(g+1)$ and $t_2(\kappa, g) = 48(\kappa+1)(g+1)$. Apply the following routine on G:

1. Let $G' := G$.
2. While $\mathbf{tw}(G') \geq t_2(\kappa, g)$ do
3. let $i = \kappa + 2$,
4. let R' be a subdivided wall of height i, as in the above claim, and
5. update G' to the subgraph of G' induced by the vertices in the *strict* interior of the perimeter of R'.
6. Output G'.

Notice that the output of the above routine has always treewidth at most $t_2(\kappa, g)$. If the above algorithm never enters the loop of lines 3–5, then $\mathbf{tw}(G') = \mathbf{tw}(G) \geq t_1(\kappa, g)$ and, because of the above claim for $i = k$, G contains the desired subdivided wall R of height k. If this is not the case, then because of the stripping of Line 5, G' (and thus G as well) contains a wall R of height $i - 2 = k$, as required. $\square$

The third assertion of Lemma 1 is important for our algorithm, as it implies that all subgraphs of G that are inside the outer cycle have bounded treewith and therefore, for these graphs, it is possible to answer queries on (rooted) surface topological minor containment in $O_h(n)$ steps.

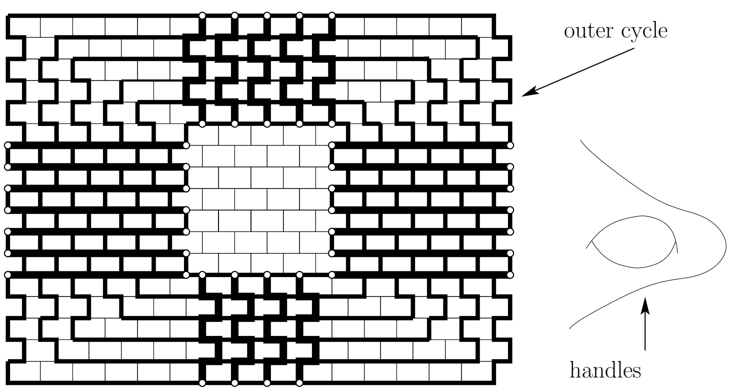

Fig. 1. A wall of height 17 and width 15 together with a railed annulus of 6 cycles and 23 rails in it.

Cycles, rails, and tracks. Notice now that inside the perimeter of a subdivided wall of "big" enough height and width, one may distinguish a collection of nested cycles $\mathcal{A} = \{C_1, \ldots, C_r\}$ all met by a collection of paths $\mathcal{W} = \{W_1, \ldots, W_q\}$ (we call them *rails*) in a way that the intersection of a rail and a cycle is always a path. We can also assume that, among these cycles, C_r is the perimeter of the subdivided wall and we call it the *outer cycle*.

See Figure 1 for an example of how to extract 6 cycles and 23 rails from a (subdivided) wall of height 17 and width 15. We call this pair $(\mathcal{A}, \mathcal{W})$ of collections of cycles and rails *railed annulus* and observe that all rails and cycles are contained inside the outer cycle. Moreover, given that we need k_1 cycles and k_2 rails, we can always find them in a subdivided wall of big enough height and width. Combining this fact with Lemma 1, we derive the following.

Lemma 2. *There exist functions t_3 and t_4, such that every graph G that is embedded in a surface Σ of Euler genus g and has treewidth at least $t_3(r, q)$ contains a* railed annulus *$(\mathcal{A}, \mathcal{W})$ if r cycles and q rails such that every subgraph of G that is entirely inside the outer cycle of $\mathcal{A}$ has treewidth at most $t_4(r, q)$.*

For a more abstract visualization of a railed annulus with 9 cycles and 24 rails, see Figure 2.

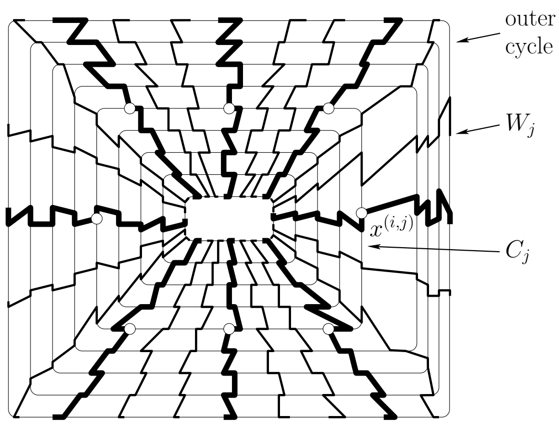

Fig. 2. A railed annulus of 9 cycles and 24 rails. Among them, we distinguish 8 tracks.

For the purposes of our algorithm, we distinguish some proper subset of the rails and we call them *tracks*. For each cycle C_i of a railed annulus and for each rail W_h, we denote by $x^{(i,j)}$ the last vertex, starting from inside, of W_h that is a vertex of C_i. For the i-th cycle (counting from inside to outside) we denote by $X^{(i)}$ the set of all $x^{(i,j)}$'s on it (in Figure 2, $X^{(5)}$ consists of the white vertices). Also, for each i, we denote by $\Delta^{(i)}$ the inner closed disk bounded by C_i and by $G^{(i)}$ the subgraph of G that is is inside $\Delta^{(i)}$.

Crossings of a pattern graph. Let H be a Σ-embedded pattern graph of at most h edges and let $\tilde{\Delta}$ be a closed disk of Σ. The notion of a graph J that is $\tilde{\Delta}$-excised by H is visualized in Figure 5. Notice that J is embedded inside $\tilde{\Delta}$ and contains new vertices (the white vertices, denoted by X) that are the points of intersection of H with the boundary of $\tilde{\Delta}$. The number of these white vertices is the *crossing number* of J. We see each $\tilde{\Delta}$-excised graph J as being embedded inside the disk Δ. We also consider its *enhancement* $J_{\tilde{\Delta},X}$ by adding edges between boundary vertices as depicted in Figure 5.

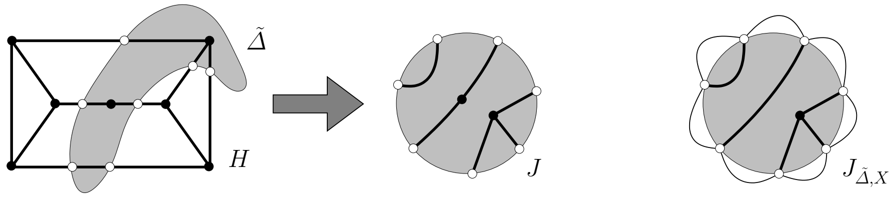

Fig. 3. A graph J that is $\tilde{\Delta}$-excised by H and its enhanced version $J_{\tilde{\Delta},X}$ (X consists of the white vertices).

We say tha two $\tilde{\Delta}$-excised graphs J^1 and J^2 are *equivalent* if their enhancements $J^1_{\tilde{\Delta},X}$ and $J^2_{\tilde{\Delta},X}$ are topologically isomorphic.

We also define the same enhancement for each graph $G^{(i)}$ and we denote it by $G^{(i)}_{\Delta^{(i)},X^{(i)}}$ (see the left part of Figure 5).

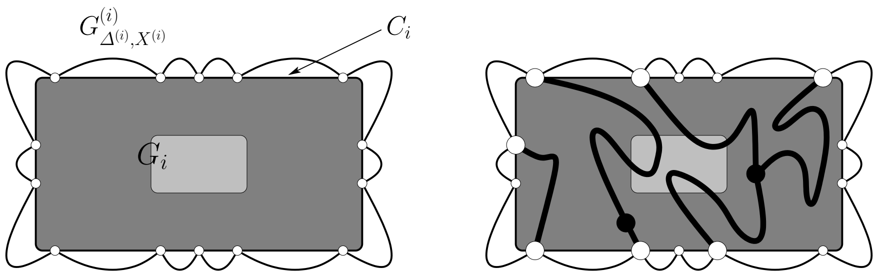

Fig. 4. The graph $G^{(i)}_{\Delta^{(i)},X^{(i)}}$ and a realization of J as a ρ-attached topological minor of $G^{(i)}$.

Attached topological minors. We set up a repository $\mathbf{H}_h$ of all graphs J that can be $\tilde{\Delta}$-excised by H with crossing number $f_4(h)$ where f_4 is a function to be determined later. Clearly, the size of $\mathbf{H}_h$ depends exclusively on h. Our next step is to set up a 0/1-vector χ_i that encodes, for every $J \in \mathbf{H}_h$ and every mapping $\rho : X \to X^{(i)}$, whether $J_{\tilde{\Delta},X}$ is a surface topological minor of $G^{(i)}_{\Delta^{(i)},X^{(i)}}$, where the vertices of X are mapped to vertices of $X^{(i)}$ as indicated by ρ. When this happens, we say that J is a *ρ-attached topological minor of* $G^{(i)}$). For an example of such a mapping, see the right part of Figure 5.

Detecting an irrelevant edge. As each $G^{(i)}$ has bounded treewidth and the property of being a ρ-attached topological minor can be expressed in MSOL, χ_i can be computed in $O_h(n)$ steps and can be encoded in space that depends exclusively on h. It is important to notice that the vector sequence $\chi_1, \ldots, \chi_r$ is *monotone* in the sense that if a graph J is a ρ-attached topological minor of G^i, then it is also a ρ-attached topological minor of $G^{i'}$ for $i' > i$. By a pigeonhole argument, if the number of the cycles in the railed annulus is big enough, then there should exist a sub-collection $C_{\theta+1}, \ldots, C_{\theta+l}$ of consecutive cycles where $\chi_{\theta+1} = \ldots = \chi_{\theta+l}$, i.e., where the members of $\mathbf{H}_r$ behave the same as ρ-attached topological minors in their interiors (here l will be chosen to be as big as required for the correctness of our proofs). We call the seqence $C_{\theta+1},\dots,C_{\theta+l}$ *frozen* and observe that it can be detected algorithmically in $O_h(n)$ steps. In other words, we have the following:

Lemma 3. *There exists some function $g:\mathbb{N}\to\mathbb{N}$ such that for every two positive integers h and l, every Σ-embedded graph G with a (r,q)-*railed annulus *$(\mathcal{A},\mathcal{W})$ where $r\geq g(h)\cdot l$, and every $I\subset\{1,\dots,q\}$ there is an integer $\theta\in\{0,\dots,r-l\}$, such that the sequence $\{\chi_1,\dots,\chi_r\}$ contains a subsequence $\{\chi_{\theta+1},\dots,\chi_{\theta+l}\}$ of l consecutive equal vectors. Moreover, there is an algorithm that, given h,l, G, $(\mathcal{A},\mathcal{W})$, and I, outputs θ in $\phi(h,\mathbf{tw}(G^{(r)}))\cdot n$ steps, for some function ϕ.*

We claim that any edge in a non-track rail that lies between C_r and C_{r+1} is an irrelevant edge. In other words, the procedure Procedure $\mathsf{Irrelevant\ Edge\ Detection}(G,\Sigma)$ is the following:

Procedure $\mathsf{Irrelevant\ Edge\ Detection}(G,\Sigma)$
1. Compute $\mathbf{H}_h$.
2. Find, using Lemma 2, a railed annulus $(\mathcal{A},\mathcal{W})$
   in G with $r=g(h)\cdot t_3(h)$ cycles and $t_4(h)$ rails.
3. Pick a proper subset I of $\{1,\dots,q\}$ of size $t_5(h)$
   and call the rails in $\{W_i\mid i\in I\}$ *tracks*.
4. Apply Lemma 3, using $(\mathcal{A},\mathcal{W})$ and its tracks, in order to
   detect a frozen sequence $C_{\theta+1},\dots,C_{\theta+l}$ in $\mathcal{A}$.
6. Let $i\in\{1,\dots.r\}\setminus I$ and let e be an edge of W_i that lies
   between $C_{\theta+1}$ and $C_{\theta+2}$, i.e. an edge in $W_i\cap(\Delta_{\theta+2}\setminus C_{\theta+1}\setminus\Delta_{\theta+2})$.
7. Output e.

The functions t_3,t_4, and t_5 above, depend on H and the genus of G and will be determined later so that the algorithm is correct.

5 Correctness of the algorithm

This section contains a sketch of the proof that irrelevant edges are indeed irrelevant.

Linkage extraction. Suppose that H is a surface topological minor of G. Our purpose is to find a realization of H as a surface topological minor of G in a way that avoids the irrelevant edge. For this we fix our attention in the "frozen" annulus defined by the cycles $C_{\theta+1}$ and $C_{\theta+l}$. As H has at most $2\cdot h$ vertices, there should be a big enough sub-annulus that does not contain any images of the vertices of H. Assume that this sub-annulus contains the r' cycles $C_{\theta+\theta'+1},\dots,C_{\theta+\theta'+r'}$. Notice that H defines a collection of disjoint paths whose terminals are outside this annulus. This collection is a *h'-linkage* (i.e. a subgraph consisting of a collection of at most h' disjoint paths) for some $h'\leq h$ and we denote it by $\mathcal{L}'$ (see Figure 5).

Linkage replacement. The terminals of a linkage are the endpoints of its paths. Recall that the terminals of the linkage $\mathcal{L}$ that we detected in the previous paragraph has all its linkages outside the closed annulus defined by the cycles C_1 and C_r. We call such a linkage $\mathcal{A}$-avoiding linkage. Our next step is to prove the following lemma:

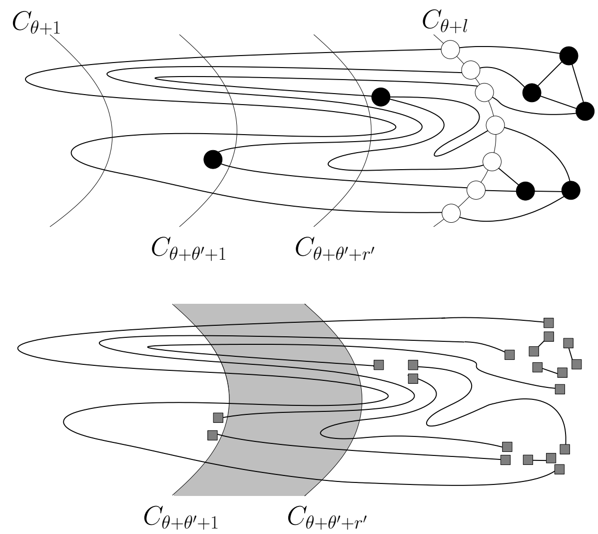

Fig. 5. The upper figure depicts a realization of H as a topological minor of G. The annulus defined by the cycles $C_{\theta+\theta'+1}$ and $C_{\theta+\theta'+r'}$ does not contain any image of a vertex in H. The lower figure shows the corresponding linkage.

Lemma 4. *There exist functions t_3, t_4, and t_5 such that the following hold: If h is a positive integer h, G a Sigma-embedded graph with a railed annulus $(\mathcal{A}, \mathcal{W})$ with $r = t_3(h)$ cycles and $q = t_4(h)$ rails, $\mathcal{L}$ an $\mathcal{A}$-avoiding linkage $\mathcal{L}$ and subset I a proper subset of $\{1, \ldots, q\}$ where $|I| = t_5(h)$, then there is an $\mathcal{A}$-avoiding linkage $\mathcal{L}$ with the following properties:*

- *the paths of $\mathcal{L}$ link the same terminals as the paths in $\mathcal{L}'$,*
- *no more than $t_5(h)$ paths in $\mathcal{L}'$ cross the "middle" cycle $C_{\lceil r/2 \rceil}$ and, when this happens, their intersection will be just a path,*
- *when we orient such a path from inside to outside, its last in C_μ should always be a vertex of $X^{(\mu)}$.*

The proof of the above lemma is quite technical and uses the "vital linkage" Theorem of Roberstong and Seymour in [25] (actually the function t_5 is directly taken from [25]). An example of this linkage replacement is depicted in Figure 6.

Pattern displacement. Our next step is to observe that the new linkage gives rise to a graph J' of $\mathbf{H}_h$ that is a ρ-attached topological minor of $G^{(\mu)}$. Recall that $\chi_{\theta+\theta'+1} = \chi_\mu$. Therefore, J' is also ρ'-attached topological minor of $G^{(\theta+\theta'+1)}$ where ρ' is the "left-side displacement" of ρ from C_μ to $C_{\theta+\theta'+1}$. But then, we may use the segments of the tracks that are cropped by the annulus defined by C_μ and $C_{\theta+\theta'+1}$ to realize J' as a ρ'-attached topological minor of $G^{(\mu)}$ in a way that rails that are not tracks are avoided (see Figure 7).

Clearly, the new realization of J' avoids the irrelevant edge and can be extended to a realization of H as a surface topological minor of G (see the right part of Figure 7). This means that that the irrelevant edge is indeed irrelevant and concludes with the proof of the correctness of procedure $\mathsf{Irrelevant\ Edge\ Detection}(G, \Sigma)$.

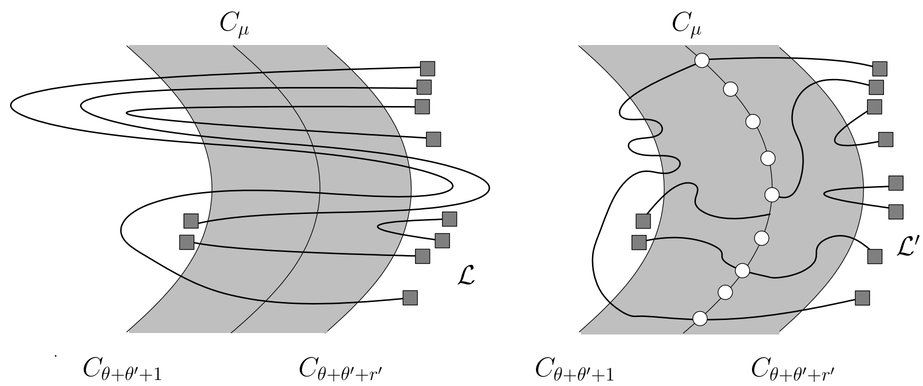

Fig. 6. The replacement of linkage $\mathcal{L}$ by a linkage $\mathcal{L}'$. (We do not depict paths that are entirely outside the sub-annulus. Also, for reasons of simplicity we represent the intersection of all, except from one, paths with C_μ by a single vertex instead of a path.)

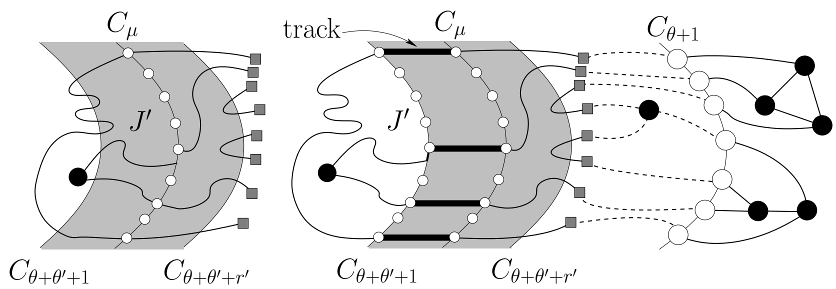

Fig. 7. Two different realizations of J' as ρ-attached topological minors of G^μ. The one on the right avoids the irrelevant edge.

6 (More) definitions

General. Given a graph G we denote by $\mathcal{C}(G)$ the set of its connected components. A *non-trivial connected component* is one that has at least two vertices. A *noose* of a Σ-embedded graph G is a simple closed curve on Σ that intersects G only at vertices and every face at most once. If $S \subset \mathbb{R}^2$ is a set of points, then $\mathbf{bor}(S)$ is the boundary of S and $\overline{S}$ is the closure of S.

Rooted graphs. Given a graph G and a subset S of its vertices, we say that the pair (G, S) is the graph G *rooted* on the set S, i.e., the set of its *roots*.

Paths. Given a graph G an *extended path* in G is a pair (P, A) such that A is a set of one or two vertices of G and P is either a path between a and b, if $|A| = 2$, or a cycle meeting A, if $|A| = 1$. The internal vertices of an extended path (P, A) are the vertices in $V(P) \setminus A$. Also, given an extended path $\hat{P} = (P, A)$ we define its vertex and edge set as $V(\hat{P}) = V(P)$ and $E(\hat{P}) = E(P)$ respectively. We say that two extended paths are *internally disjoint* if no internal vertex of one of them belongs to the vertex set the other. If $\mathcal{P}$ is a collection of extended paths in a graph H, we denote $V(\mathcal{P}) = \bigcup_{\hat{P} \in \mathcal{P}} V(\hat{P})$, $E(\mathcal{P}) = \bigcup_{\hat{P} \in \mathcal{P}} E(\hat{P})$ and $G[\mathcal{P}] = (V(\mathcal{P}), E(\mathcal{P}))$.

7 Linkages through an annulus

Let G be a Σ-embedded graph. An r-*linkage* in G is a set of r pairwise disjoint paths of it. The *endpoints* of a linkage $\mathcal{L}$ are the endpoints of the paths in $\mathcal{L}$. The *pattern* of $\mathcal{L}$ is defined as

$$\pi(\mathcal{L}) = \{\{s,t\} \mid \mathcal{L} \text{ contains a path from } s \text{ to } t\}.$$

Two linkages $\mathcal{L}, \mathcal{L}'$ are *equivalent* if they have the same pattern, i.e. $\pi(\mathcal{L}) = \pi(\mathcal{L}')$. A linkage is *vital* if it spans all vertices and no other linkage connects the same pairs of end vertices.

Proposition 3 ((1.1) in [25]). *There exists a function $f_1 : \mathbb{N} \to \mathbb{N}$ such that every graph with a vital k-linkage has treewidth at most $f_1(k)$.*

Let C_{out} be a contractible cycle in G and let Δ_{out} be a connected component of the set $\Sigma \setminus C_{\text{out}}$ such that the interior of $\Sigma \setminus \Delta_{\text{out}}$ is homeomorphic to an open disk. Let also C_{in} be a cycle in the interior of $\Sigma \setminus \Delta_{\text{out}}$ and let Δ_{in} be the connected component of the set $\Sigma \setminus C_{\text{in}}$ that does not contain C_{out}. We say that a closed subset Ω of Σ is an *annulus* if there are disks Δ_{out} and Δ_{in} as before, such that $\Omega = \Sigma \setminus (\Delta_{\text{in}} \cup \Delta_{\text{out}})$. We call the open set Δ_{in} the *interior* of Ω and we call the open set Δ_{out} the *exterior* of Ω. Notice that if Σ is the sphere, Δ_{in} and Δ_{out} can be interchanged while this is not the case for other surfaces.

Let Ω be an annulus of Σ. We say that a linkage $\mathcal{L}$ is Ω-*avoiding* if none of the terminals of $\mathcal{L}$ lies inside Ω. Given that the linkage $\mathcal{L}$ is Ω-avoiding, we define

$$\mathcal{K}(\mathcal{L}, \Omega) = \bigcup_{P \in \mathcal{L}} \{Z \mid Z \text{ is a connected component of the set } P \cap \Omega\}$$

and observe that each member of $\mathcal{K}(\mathcal{L}, \Omega)$ is a path in G. Therefore, $\mathcal{K}(\mathcal{L}, \Omega)$ is a linkage of G with all its endpoints lying on the boundary of Ω.

An r-*annulus* of a Σ-embedded graph G is a collection $\mathcal{A} = \{C_1, \ldots, C_r\}$ of r mutually vertex disjoint contractible cycles of G with the property that there exists an annulus Ω in Σ where

- all cycles in $\mathcal{A}$ are inside Ω,
- C_1 (resp. C_r) is the boundary of the interior (resp. exterior) of Ω, and
- for $i = 2, \ldots, r-1$, one of the connected components of the set $\Sigma \setminus C_i$ contains all cycles in $\{C_1, \ldots, C_{i-1}\}$ and the other contains all cycles in $\{C_{i+1}, \ldots C_r\}$.

We call Ω the *territory* of $\mathcal{A}$ and we denote it by $\Omega_{\mathcal{A}}$.

Given two integers i, j where $1 \leq i < j \leq r$ we set $\mathcal{A}[i,j] = \{C_i, \ldots, C_j\}$ and notice that $\mathcal{A}[i,j]$ is an $(i-j+1)$-annulus. We define $G_{\mathcal{A}} = \bigcup_{C \in \mathcal{A}} C$ and we denote by $\Delta^{(i)}$ the closed disk bounded by C_i that contains the interior of $\Omega_{\mathcal{A}}$.

We say that a linkage $\mathcal{L}$ is $\mathcal{A}$-*avoiding* when it is $\Omega_{\mathcal{A}}$-avoiding. Let $\mathcal{A} = \{C_1, \ldots, C_r\}$ be an r-annulus of a Σ-embedded graph G and let $\mathcal{L}$ be an $\mathcal{A}$-avoiding linkage of G. We call a path in $\mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$ $\mathcal{A}$-*crossing* if one of its endpoints is in C_1 and the other is in C_r, otherwise we call it *non-$\mathcal{A}$-crossing*. Let P be a non-$\mathcal{A}$-crossing path of $\mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$ and assume that its endpoints belong both to C, where C is either C_1 or C_r. Notice that, in case $P \not\subseteq C$, $C \cup P$ defines three closed disks: among them, we denote by Δ_P be the only one that is inside $\Omega_{\mathcal{A}}$. If $P \subseteq C$, we set $\Delta_P = P$.

We say that a path $P \in \mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$ is *maximal* if there is no other $P' \in \mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$ such that $\Delta_P \subset \Delta_{P'}$. The *depth*, denoted $\mathbf{depth}_{\mathcal{L},\mathcal{A}}(P)$, of the non-crossing path P in $\mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$ whose both endpoints lie in C_1 (resp. C_r) is the maximum i such that P intersects all cycles in $\{C_1, \ldots, C_i\}$ (resp. $\{C_{r-i+1}, \ldots, C_r\}$). For completeness, we define the depth of a crossing path to be zero.

Let G be a Σ-embedded graph and let $\mathcal{A}$ be an annulus of G. Given an $\mathcal{A}$-avoiding k-linkage $\mathcal{L}$, we define

$$c(\mathcal{L}) = \sum_{R \in \mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})} \mathbf{depth}_{\mathcal{L},\mathcal{A}}(R) + |\bigcup_{P \in \mathcal{L}} (E(P) \setminus E(G_{\mathcal{A}}))|.$$

Let D be a closed disk. We say that an $\mathcal{A}$-avoiding k-linkage $\mathcal{L}$ is $(\mathcal{A}, D)$-*minimal* if among all linkages $\mathcal{L}'$ that are equivalent to $\mathcal{L}$ and such that $V(\mathcal{L}) \cap D = \emptyset$, $c(\mathcal{L})$ is minimized.

Lemma 5. *Let k, z be two positive integers and let f_1 be the function in Proposition 3. Let also G be a Σ-embedded graph, $\mathcal{A}$ be a y-annulus of G where $y = z + 4 \cdot f_1(k)$, D be a closed disk where $D \subseteq \Omega_{\mathcal{A}}$, and $\mathcal{L}$ be a $(\mathcal{A}, D)$-minimal $\mathcal{A}$-avoiding k-linkage of G. Then,*

1. *none of the non-$\mathcal{A}$-crossing paths of $\mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$ meets the territory of $\mathcal{A}[4 \cdot f_1(k), z + 4 \cdot f_1(k)]$,*
2. *$\mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$ has at most $f_1(k)$ $\mathcal{A}$-crossing paths.*

Proof. For our proof we need first some definitions and observations. We first define

$$H = G[(\bigcup_{i=1,\ldots,y} C_i) \cup (\bigcup_{P \in \mathcal{L}} P)]$$

and for each path $R \in \mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$ we define $\mathcal{Q}_R$ to be the set of non-trivial connected components of $R \cap G_{\mathcal{A}}$, i.e. $\mathcal{Q}_R$ contains the subpaths of R that have common edges with the cycles of $\mathcal{A}$. We now define $\hat{H}$ as the graph obtained from H if, for each $Q \in \bigcup_{R \in \mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})} \mathcal{Q}_R$ we contract all edges of Q to a single vertex v_Q. Let also $E^+ \subseteq E(H)$ be the set of all contracted edges. Notice that all non-terminal vertices of $\hat{H}$ have degree either 2 or 4. Similarly, we define $\hat{\mathcal{A}} = \{\hat{C}_1, \ldots, \hat{C}_y\}$ by contracting, in each C_i, its common edges with E^+. As all edges of E^+ are edges of the paths in $\mathcal{L}$, it follows that their contraction can transform $\mathcal{L}$ to a linkage $\hat{\mathcal{L}}$ of $\hat{H}$ with the same pattern as $\mathcal{L}$. For each path $P \in \mathcal{L}$ we use the notation $\hat{P}$ for its counterpart in $\hat{\mathcal{L}}$.

Observe also that any linkage $\hat{\mathcal{L}}^\star$ of $\hat{H}$ can be turned back to a linkage $\mathcal{L}^\star$ of H with the same pattern as $\hat{\mathcal{L}}^\star$ by uncontracting each vertex v_Q of each path of $\hat{\mathcal{L}}^\star$ to the path Q. Notice that $\hat{\mathcal{L}}$ is an $\hat{\mathcal{A}}$-avoiding k-linkage in $\hat{H}$. Finally, by the definition of E^+, it follows that for every linkage $\mathcal{L}^\star$ of H,

$$\forall \hat{R} \in \mathcal{K}(\hat{\mathcal{L}}^\star, \Omega_{\hat{\mathcal{A}}}) \ \ \mathbf{depth}_{\hat{\mathcal{L}}^\star,\hat{\mathcal{A}}}(\hat{R}) = \mathbf{depth}_{\mathcal{L}^\star,\mathcal{A}}(R) \tag{1}$$

$$|\bigcup_{\hat{P} \in \hat{\mathcal{L}}^\star} E(\hat{P}) \setminus E(G_{\hat{\mathcal{A}}})| = |\bigcup_{P \in \mathcal{L}^\star} E(P) \setminus E(G_{\mathcal{A}})| \tag{2}$$

Our target is to prove that $\mathcal{L}$ satisfies the two conditions of the lemma. Let E' be the edges of H inside D. Observe that these edges remain intact in $\hat{H}$. Let $\hat{H}^-$ be $\hat{H}$ without the edges of E'. As the first step we claim that violation of one of them implies that $\mathbf{tw}(\hat{H}^-) > f_1(k)$.

Suppose that the first condition of the lemma is violated and there exists a maximal non-$\mathcal{A}$-crossing path R_h in $\mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$ with depth at least $h = 4 \cdot f_1(k)+1$. Without loos of generality, we assume that the endpoints of R_h are in C_1. This means that R_h meets C_h, which implies that $\hat{R}_h$ meets $\hat{C}_h$. We first claim that there is a path $\hat{R}_{h-1} \in \mathcal{K}(\hat{\mathcal{L}}, \Omega_{\hat{\mathcal{A}}})$ (that is also a path in $\hat{H}$) meeting $\hat{C}_{h-1}$. Indeed, if this is not the case, it is possible to replace some subpath of $\hat{R}_h$ by some subpath of $\hat{C}_{h-1}$ and construct a new path $\hat{R}^\star$ such that $\hat{\mathcal{L}}^\star = \hat{\mathcal{L}} \setminus \{\hat{R}\} \cup \{\hat{R}^\star\}$ is an other linkage of in $\hat{H}$ that is equivalent to $\hat{\mathcal{L}}$. By the construction of $\hat{R}^\star$, it follows that $\mathbf{depth}_{\hat{\mathcal{L}}^\star,\hat{\mathcal{A}}}(\hat{R}^\star_h) < \mathbf{depth}_{\hat{\mathcal{L}},\hat{\mathcal{A}}}(\hat{R}_h)$. Let $R^\star_h$ be the counterpart of $\hat{R}^\star_h$ in H. As we mentioned before, $\hat{\mathcal{L}}^\star$ corresponds to a linkage $\mathcal{L}^\star$ in H with the same pattern and, from (1), it follows that $\mathbf{depth}_{\mathcal{L}^\star,\mathcal{A}}(R^\star_h) = \mathbf{depth}_{\hat{\mathcal{L}}^\star,\hat{\mathcal{A}}}(\hat{R}^\star_h) < \mathbf{depth}_{\hat{\mathcal{L}},\hat{\mathcal{A}}}(\hat{R}_h) = \mathbf{depth}_{\mathcal{L},\mathcal{A}}(\hat{R}_h)$. This, in turn, implies that $c(\mathcal{L}^\star) < c(\mathcal{L})$, a contradiction to the choice of $\mathcal{L}$ and the claim holds. Using the same argument, we conclude that, for every $i = h, \ldots, 1$, there should be a path $\hat{R}_i \in \mathcal{K}(\hat{\mathcal{L}}, \Omega_{\hat{\mathcal{A}}})$ in $\hat{H}$ meeting all paths in $\{\hat{C}_i, \ldots, \hat{C}_1\}$. This implies that $\hat{H}$ contains an $(h/2 \times h/2)$-grid as a minor created by the intersection of $C_{h/2+1}, \ldots, C_{h+1}$ with $\hat{R}_1, \ldots, \hat{R}_{h/2}$. Consequently, $\hat{H}$ contains two disjoint $(h/4 \times h/4)$-grids Γ_1 and Γ_2 as a minor. Γ_1 is created by the intersection of $C_1, \ldots, C_{h/4}$ and $\hat{R}_{h/2}, \ldots, \hat{R}_{3h/4}$ and Γ_2 is created by the intersection of $C_{h/4+1}, \ldots, C_{h/2}$ and $\hat{R}_{3h/4+1}, \ldots, \hat{R}_h$.

Notice that the connected components of the set $\Omega_{\hat{\mathcal{A}}} \setminus (C_1 \cup C_y \cup \bigcup_{j=1,\ldots,h} \hat{R}_j)$ are open disks $D_1, \ldots, D_{h+1}$, assuming that R_i is contained in the boundary of D_i and D_{i+1}, for $i = 1, \ldots, h$. Clearly, exactly one of them, say D_m, contains D. If $m > h/2$, then the model of Γ_1 avoids the vertices in D; otherwise, the model of Γ_2 avoids the vertices in D. Therefore, in both cases, $\mathbf{tw}(\hat{H}^-) \geq h/4 > f_1(k)$

Suppose now that the second condition of the lemma is violated. Then there are more than $f_1(k)$ crossing paths in $\mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$. This implies that there are more than $f_1(k)$ crossing paths in $\mathcal{K}(\hat{\mathcal{L}}, \Omega_{\hat{\mathcal{A}}})$ as well. These paths, together with the $y > f_1(k)$ cycles of $\hat{\mathcal{A}}$ imply that $\hat{H}$ contains an $((f_1(k)+1) \times (f_1(k)+1))$-grid as a minor, therefore $\mathbf{tw}(\hat{H}) > f_1(k)$.

Applying Proposition 3 we obtain that $\hat{\mathcal{L}}$ is not a vital linkage of $\hat{H}^-$, therefore there is another equivalent one, say $\hat{\mathcal{L}}^\star$, in $\hat{H}^-$, which is an $\mathcal{A}$-avoiding k-linkage of G such that $V(\mathcal{L}^\star) \cap D = \emptyset$. Let $\hat{P}^\star_{a,b}$ be an (a,b)-path in $\hat{\mathcal{L}}^\star$ that does not exist in $\hat{\mathcal{L}}$ and let $\hat{P}_{a,b}$ be the corresponding (a,b)-path of $\hat{\mathcal{L}}$. Let also e be the first edge of $\hat{P}_{a,b}$, starting from a, that is not an edge of $\hat{P}^\star_{a,b}$. Clearly, e is an edge that lies in the territory of $\hat{\mathcal{A}}$, one of its endpoints (the one closer to a) has degree 4 (otherwise e would not be the first such edge) and therefore is not used by any other path in $\hat{\mathcal{L}}^\star$. From (2), it follows that $|\bigcup_{P\in\mathcal{L}} E(P) \setminus E(G_{\mathcal{A}})| = |\bigcup_{\hat{P}\in\hat{\mathcal{L}}} E(\hat{P}) \setminus E(G_{\hat{\mathcal{A}}})| = |\bigcup_{\hat{P}\in\hat{\mathcal{L}}} E(\hat{P})| > |\bigcup_{\hat{P}^\star\in\hat{\mathcal{L}}^\star} E(\hat{P}^\star) \setminus E(G_{\hat{\mathcal{A}}^\star})| = |\bigcup_{P\in\mathcal{L}^\star} E(P) \setminus E(G_{\mathcal{A}})|$ which, in turn implies that $c(\mathcal{L}^\star) < c(\mathcal{L})$, a contradiction with the minimality of $\mathcal{L}$. $\square$

Let $\mathcal{A}$ be an annulus of a graph G. Two vertices of G are $\mathcal{A}$-*homocyclic* if they belong to the same cycle of $\mathcal{A}$. Let $\mathcal{L}$ be a linkage of G. Two vertices of G are $(\mathcal{A}, \mathcal{L})$-*colinear* if they are belong to the same path of $\mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$.

Let v, w be two $\mathcal{A}$-homocyclic and $(\mathcal{A}, \mathcal{L})$-colinear vertices of G. Let C_i be the cycle of $\mathcal{A}$ containing v and w and P be the path of $\mathcal{K}(\mathcal{L}, \Omega_{\mathcal{A}})$ containing v and w. Let also $P_{v,w}$ be the subpath of P with endpoints v and w. We define the $(\mathcal{A}, \mathcal{L})$-*amplitude* of the pair v, w to be $\max\{|i-j| \mid \exists\, j : V(C_j) \cap V(P_{v,w}) \neq \emptyset\}$.

Lemma 6. *Let f_1 be the function of Proposition 3. Let G be a Σ-embedded graph, $\mathcal{A}$ be an annulus of G, and D be a closed disk where $D \subseteq \Omega_{\mathcal{A}}$. Let $\mathcal{L}$ be a $(\mathcal{A}, D)$-minimal $\mathcal{A}$-avoiding k-linkage. Then, the $(\mathcal{A}, \mathcal{L})$-amplitude of any two $\mathcal{A}$-homocyclic and $(\mathcal{A}, \mathcal{L})$-colinear vertices is at most $2 \cdot f_1(k)$.*

Proof. We define H, $\hat{H}$, and $\hat{H}^-$ as in the proof of Lemma 5. Let v and w be a pair of two $\mathcal{A}$-homocyclic with respect to a cycle C_i and $(\mathcal{A}, \mathcal{L})$-colinear vertices such that their distance along C_i is minimal. Suppose for a contradiction that $P_{v,w}$ intersects a cycle C_j such that $|i - j| > 2 \cdot f_1(k)$. Without loss of generality, we assume that $j > i$. As in the proof of Lemma 5, we set $h = 2 \cdot f_1(k) + 1$ and we find a collection of nested paths $\hat{R}_1, \ldots, \hat{R}_h$ in $\hat{H}$, whose endpoints are contracted subpaths of C_i. (Here, $\hat{R}_h$ is the contracted counterpart of $P_{v,w} = R_h$.) Again, by the same arguments as in the proof of Lemma 5, these paths together with the contracted cycles $\hat{C}_i, \ldots, \hat{C}_j$ certify that the $\mathbf{tw}(\hat{H}^-) > 2 \cdot f_1(k)$. Using the argument from the last paragraph of the proof of Lemma 5, we have a contradiction with the minimality of $\mathcal{L}$. $\square$

Let G be a Σ-embedded graph and let $\mathcal{A} = \{C_1, \ldots, C_r\}$ be an r-annulus in G. Suppose now that, additionally, $\mathcal{W} = \{W_1, \ldots, W_q\}$ is a collection of mutually vertex-disjoint paths of $G \cap \Omega_{\mathcal{A}}$ such that for each $(i, j) \in [r] \times [q]$ $C_i \cap W_j$ is a path $P^{(i,j)}$ of G. Then we call the pair $(\mathcal{A}, \mathcal{W})$, an (r, q)-*railed annulus* of G and the paths in $\mathcal{W}$ *rails* of it. We see each W_j as being oriented from the interior towards the exterior of $\Omega_{\mathcal{A}}$. Having this orientation in mind, we denote by $x^{(i,j)}$ the finishing endpoint of $P^{(i,j)}$. Let $J \subseteq \{1, \ldots, q\}$ and set $X_J^{(i)} = \{x^{(i,j)} \mid j \in J\}$. Given $(i, h) \in [r]^2$, we define a bijection $\lambda_{i,h}^J : X_J^{(i)} \to X_J^{(h)}$ letting $\lambda_{i,h}^J(x^{(i,j)}) = x^{(h,j)}$ for all $j \in J$. When J is clear from the context, we use notation $\lambda_{i,h}$ instead of $\lambda_{i,h}^J$.

Given a path P of a Σ-embedded graph, a closed disk whose boundary is a noose and a vertex x of P in this noose, we say that P *abandons* Δ *in* x if x is the unique vertex in $V(P) \cap \Delta$ that is adjacent to a vertex in $P \setminus \Delta$. We also say that a linkage $\mathcal{L}$ *abadons* Δ *on* X if every $L \in \mathcal{L}$ abandons Δ in some vertex of X.

We now give the main result of this section that is is proved using Lemma 5.

Lemma 7. *Let k be a positive integer. Let also G be a Σ-embedded graph, $(\mathcal{A}, \mathcal{W})$ be a (r, q)-railed annulus of G where $r = 4 \cdot (f_1(k))^2 + 6 \cdot f_1(k) + 1$ and $q = 5 \cdot f_1(k)$, and let $\mathcal{L}'$ be an $\mathcal{A}$-avoiding k-linkage of G. Then for each subset $I \subseteq \{1, \ldots, q\}$ where $|I| = f_1(k)$ there is an $\mathcal{A}$-avoiding k-linkage $\mathcal{L}$ that is equivalent to $\mathcal{L}'$ and such that $\mathcal{L}'$ abandons $\Delta^{(y)}$ on $X_I^{(y)}$ where $y = (r + 1)/2$.*

Proof. Let $\mathcal{A}^*$ be a collection of cycles $C_1^*, \ldots, C_{q/2}^*$ such that C_i^* is the unique cycle of the graph $W_i \cup W_{q-i} \cup C_i \cup C_{r-1}$ that intersects the rails $W_i, \ldots, W_{q-i}$ but not any other rail, for $i = 1, \ldots, q/2$. Actually, we redefine $C_{q/2}^*$ to be a face inside the closed disk bounded by $C_{q/2}^*$ and does not contain any other cycle from $C_1^*, \ldots, C_{q/2-1}^*$.

Let D^* be any closed disk in $\Omega_{\mathcal{A}^*}$ that does not contain any vertex. Let $\mathcal{L}''$ be an $(\mathcal{A}^*, D^*)$-minimal $\mathcal{A}^*$-avoiding k-linkage that is equivalent to $\mathcal{L}'$. Clearly, $V(\mathcal{L}) \cap D^* = \emptyset$. From first condition of Lemma 5, no path from $\mathcal{K}(\mathcal{L}'', \Omega_{\mathcal{A}^*})$ meets $C_{f_1(k)}^*$.

Let D be the closed disk bounded by $C_{f_1(k)}^*$ and not containing any vertex of C_1^*. Let $\mathcal{L}'''$ be an $(\mathcal{A}, D)$-minimal $\mathcal{A}$-avoiding k-linkage that is equivalent to $\mathcal{L}'$. From Lemma 5, no non-$\mathcal{A}$-crossing path in $\mathcal{K}(\mathcal{L}''', \Omega_{\mathcal{A}})$ meets the annulus $\mathcal{A}[2f_1(k) + 1, r - 2f_1(k)]$ and there are ξ at most $f_1(k)$ $\mathcal{A}$-crossing paths in $\mathcal{K}(\mathcal{L}''', \Omega_{\mathcal{A}})$. Let $P_1, \ldots, P_\xi$ be these paths ordered with respect to the cyclic order they appear in the annulus $\mathcal{A}$. Moreover, we assume that one of the two open disks of $\Omega_{\mathcal{A}} \setminus (C_1 \cup C_r \cup P_1 \cup P_\xi)$ contains D and the closure of the other one contains $P_1, \ldots, P_\xi$.

Let Q_i^{up} be the unique path in $C_{2f_k+i(f_1(k)+1)}$ connecting $P^{(2f_k+i2f_k+1,1)}$ and $P^{(2f_k+i2f_k+1,f_1(k)-i+1)}$ that is entirely inside D, where $i = 1, \ldots, f_1(k)$. Let Q_i^{down} be the unique path in $C_{r-2f_k-i2f_1(k)}$ connecting $P^{(r-2f_k-i2f_1(k),1)}$ and $P^{(r-2f_k-i2f_1(k),f_1(k)-i+1)}$ that is entirely inside D, where $i = 1, \ldots, f_1(k)$. Let S_i be the unique path in W_i connecting $P^{(2f_1(k)+i2f_1(k)+1,f_1(k)-i+1)}$ and $P^{(r-2f_1(k)-i2f_1(k),f_1(k)-i+1)}$ that is entirely inside D, where $i = 1, \ldots, f_1(k)$.

For $i = 1, \ldots, f_1(k)$, among all the vertices that belong to both P_i and $C_{2f_1(k)+i2f_1(k)+1)}$ choose the unique vertex α_i^{up} such that there exists a subpath Z_i^{up} of $C_{2f_1(k)+i2f_1(k)+1}$ between α_1^{up} and a vertex in Q_i^{up} that do not contain any other vertex of P_i and any other vertex of Q_i^{up}. We define α_i^{down} and Z_i^{down} symmetrically. For $i = 1, \ldots, f_1(k)$, let P_i^{up} be the subpath of P_i starting from an endpoint in C_1 and finishing in α_i^{up}. We define P_i^{down} symmetrically.

We claim that for every $i = 1, \ldots, f_1(k)$, P_{i-1}^{up} does not meet Z_i, where P_0 is P_ξ. Notice that otherwise P_{i-1}^{up} would contain two vertices v, w of cycle $C_{2f_1(k)+(i-1)f_1(k)+1}$ and at least one point belonging to $C_{2f_1(k)+if_1(k)+1}$. Therefore v and w are $\mathcal{A}$-homocyclic and $(\mathcal{A}, \mathcal{L})$-colinear with $(\mathcal{A}, \mathcal{L})$-amplitude larger than $2f_1(k)$; a contradiction with Lemma 6.

Let us consider the graph $J = J_{\text{in}} \cup J_{\text{mid}} \cup J_{\text{out}}$ where

$$J_{\text{out}} = \Big(\bigcup_{P \in \mathcal{L}'''} P\Big) \cap \overline{\Sigma \setminus \Omega_{\mathcal{A}}},$$

$$J_{\text{mid}} = \bigcup_{\substack{P \in \mathcal{K}(\mathcal{L}''', \Omega_{\mathcal{A}}) \\ \setminus \{P_1, \ldots, P_\xi\}}} P, \quad \text{and}$$

$$J_{\text{in}} = \bigcup_{i=1,\ldots,\xi} P_i^{up} \cup Z_i^{up} \cup Q_i^{up} \cup S_i \cup Q_i^{down} \cup Z_i^{down} \cup P_i^{down}.$$

From the previous claim it follows that the connected components of J form a linkage $\mathcal{L}''''$ that is equivalent to $\mathcal{L}'''$. Let $\mathcal{A}^+ = \mathcal{A}[2f_1(k)+2(f_1(k))^2+1, r-(2f_1(k)+2(f_1(k))^2)]$. Observe that $\bigcup_{P \in \mathcal{L}''''} \cap \mathcal{A}^+ = \bigcup_{f_1(k)-\xi+1} W_i \cap \mathcal{A}^+$, ie. the linkage crosses $\mathcal{A}^+$ "vertically", following subpaths of its rails. This makes it possible to reroute $\mathcal{L}''''$ to a new linkage $\mathcal{L}$ with the claimed properties.

8 Looking for an irrelevant edge

Attached topological minors. Let Δ be a closed disk of Σ. We say that a graph G is Δ-*embedded* if it is embedded in Δ. Let G be such a Δ-embedded graph and let $X \subseteq V(G) \cap \mathbf{bor}(\Delta)$. We define the graph $G_{\Delta,X}$, embedded in Σ as the graph obtained by first copying the embedding of G in Σ and then adding edges (outside Δ) between any two vertices of X that lie consecutively along the boundary of the disk Δ. Let now H (G) be a $\tilde{\Delta}$-embedded (resp. Δ-embedded) graph, let $X \subseteq V(G) \cap \mathbf{bor}(\Delta)$, and consider an injection $\rho : V(H) \cap \mathbf{bor}(\tilde{\Delta}) \to X$. We say that H is a *ρ-attached topological minor* of G if $H_{\tilde{\Delta}, V(\tilde{H}) \cap \mathbf{bor}(\tilde{\Delta})}$ is a surface ρ-rooted topological minor of $G_{\Delta,X}$.

Crossings. We say that a closed curve N *crosses* an edge e of a graph if their intersection is a finite set X of points of Σ where for each $x \in X$, there are points of e in both connected components of $\Sigma \setminus N$ in any open neighborhood of x.

Let H be a graph embedded in Σ. Let also $\tilde{\Delta}$ be a closed disk of Σ whose boundary avoids the vertices of H and crosses each edge it intersects. The *crossing number* of $\tilde{\Delta}$ in H is $|\mathbf{bor}(\tilde{\Delta}) \cap H|$ i.e. the number of points of Σ that are in the intersection of the edges of G and $\mathbf{bor}(\tilde{\Delta})$. Given a closed disk $\tilde{\Delta}$ of Σ, we say that the $\tilde{\Delta}$-embedded graph J *is $\tilde{\Delta}$-excised by H* if J can be taken if we remove from H all points that are not in $\tilde{\Delta}$ and then declare as vertices the points in $\mathbf{bor}(\tilde{\Delta}) \cap H$. Given H we define:

$$\begin{aligned}\mathbf{H}_h = \{J \mid\ & J \text{ is a } \tilde{\Delta}\text{-embedded graph } \tilde{\Delta}\text{-excised by some } \Sigma\text{-embedded graph } H \\ & \text{of at most } h \text{ edges and such that the crossing number of } \tilde{\Delta} \text{ in } H \text{ is} \\ & \text{at most } f_1(h)\}.\end{aligned}$$

In the definition above, two $\tilde{\Delta}$-embeddible graphs $J^{(1)}$ and $J^{(2)}$ of $\mathbf{H}_h$ are not distinguishable if $J^{(1)}_{\tilde{\Delta},\mathbf{bor}(\tilde{\Delta})\cap V(H)}$ and $J^{(2)}_{\tilde{\Delta},\mathbf{bor}(\tilde{\Delta})\cap V(H)}$ are topologically isomorphic.

Characteristic functions. Let h be a positive integer. Let G be a Δ-embedded graph and $X \subseteq G \cap \mathbf{bor}(\Delta)$ be a set of vertices where $|X| = f_1(h)$. We define the function $\chi_{G,\,X}$ such that, given a $\tilde{\Delta}$-embedded graph $J \in \mathbf{H}_h$ and an injection $\rho : V(J) \cap \mathbf{bor}(\tilde{\Delta}) \to X$, its value is defined as follows.

$$\chi_{G,\,X}(J, \rho) = \begin{cases} 1 \text{ if } J \text{ is a } \rho\text{-attached topological minor of } G \\ 0 \text{ otherwise} \end{cases}$$

Let G be a Σ-embedded graph, let $(\mathcal{A}, \mathcal{W})$ be an (r, q)-*railed annulus* of G, and let $I \subseteq \{1, \ldots, q\}$ where $|I| = f_1(h)$. We denote by $G^{(i)}$ the $\Delta^{(i)}$-embedded graph $G \cap \Delta^{(i)}$. For $i \in \{1, \ldots, r\}$, we define $\chi_i = \chi_{G^{(i)}, X_I^{(i)}}$. We say that $\chi_i \leq \chi_h$ if for every $J \in \mathbf{H}_h$ and every injection $\rho : V(J) \cap \mathbf{bor}(\Delta) \to X_I^{(i)}$, it holds that $\chi_i(J, \rho) \leq \chi_h(J, \lambda_{i,h} \circ \rho)$.

Lemma 8. *For every $(i, h) \in [r]^2$, $i \leq j \Rightarrow \chi_i \leq \chi_h$.*

Proof. Let ρ be an injection $\rho : V(J) \cap \mathbf{bor}(\Delta) \to X_I^{(i)}$ and J be a ρ-attached topological minor of $\chi_{G^{(i)}, X_I^{(i)}}$. Notice that there are $|X_I^{(i)}|$ vertex disjoint paths in the (r, q)-railed annulus between vertices of $X_I^{(i)}$ and $X_I^{(h)}$, and they can be used to extend J be a ρ-attached topological of $\chi_{G^{(i)}, X_I^{(h)}}$. Hence, $i \leq j \Rightarrow \chi_i \leq \chi_h$, for every $(i, h) \in [r]^2$. ☐

Using Lemma 8 we can prove the following.

Lemma 9. *There exists some function $g : \mathbb{N} \to \mathbb{N}$ such that for every positive integers h, l, every Σ-embedded graph G, every (r, q)-*railed annulus *$(\mathcal{A}, \mathcal{W})$ of G where $r \geq g(h) \cdot l$, and every $I \subseteq [q]$ where $|I| = f_1(h)$, there is an integer $\theta \in \{0, \ldots, r - l\}$, such that the sequence $\{\chi_1, \ldots, \chi_r\}$ contains a subsequence $\{\chi_{\theta+1}, \ldots, \chi_{\theta+l}\}$ of l consecutive equal elements. Moreover, there is an algorithm that, given h, l, G, $(\mathcal{A}, \mathcal{W})$, and I, outputs θ in $\phi(h, \mathbf{tw}(G^{(r)})) \cdot |V(G)|$ steps, for some function ϕ.*

Proof. Notice that the size of $\mathbf{H}_h$ is bounded by a function of h. As the size of each $X_I^{(i)}, i \in \{1, \ldots, r\}$ is also bounded by $f_1(h)$ then the number of different values for χ_i

is also bounded by some function, say g, of h. As $r \geq l \cdot g(h)$, Lemma 8 implies that there will be at least l consecutive elemenets in $\{\chi_1, \ldots, \chi_r\}$. The algorithmic part of the lemma follows from the fact that computing χ_i can be reduced to a problem of checking topological minor containment that can be expressed in MSOL. ◻

Let (G, S) be a rooted graph. We say that an edge e of G is *h-irrelevant* if for every rooted graph (H, S_H) on h edges and a bijection $\sigma : S_H \to S_G$, (H, S_H) is a surface Σ-rooted topological minor of (G, S_G) if and only if (H, S_H) is a surface Σ-rooted topological minor of $(G \setminus e, S_G)$.

Calibrated subdivided walls. A *wall* of height h is the graph obtained from a $((h+1) \times (2h+2))$-grid with vertices (x, y), $x \in \{1, \ldots, 2h+4\}$, $y \in \{1, \ldots, h+1\}$ after the removal of the "vertical" edges $\{(x, y), (x, y+1)\}$ for odd $x + y$ and then the removal of all vertices of degree 1. The *perimeter* of a wall is the cycle bounding its outerface. A *subdivided wall* id the graph taken by some wall after subdividing edges.

Lemma 10. *For every positive integer κ and a surface Σ, there exist integers t and T such that in every Σ-embedded graph of treewidth at least t there exists a disk such that the graph induced by the vertices inside the disk is of treewidth at most T and contains a subdivision of a wall of height κ (we call such a wall* a calibrated subdivided wall*). Moreover, it is possible to find such a subdivided wall in quadratic time.*

For the proof of Lemma 10, we need the following result that can be deduced from Lemma 4 in [11].

Proposition 4 ([11]). *Let G be a graph embedded in a surface Σ of Euler genus g. If the treewidth of G is more than $12r(g+1)$, then G has the $(r \times r)$-grid as a surface minor. It is possible to find the grid minor in polynomial time.*

Proof (of Lemma 10). Let G be a Σ-embedded graph and g the Euler genus of Σ. First we note that every $(4i - 1 \times 2i)$-grid surface minor contains a subdivided wall of height i as a surface topological minor (for $i > 0$). From Proposition 4, if the $\mathbf{tw}(G) \geq 48i(g+1)$, then G contains a subdivided wall of height i as a surface topological minor. That also means that there exists a disk in the Σ-embedded graph G containing the subdivided wall.

Let $t = 48\kappa(g+1)$ and $T = 48(\kappa+1)(g+1)$. If $\mathbf{tw}(G) > T$, then G contains a subdivided wall of height $\kappa + 1$. Let P be the perimeter of the subdivided wall and let G' be the graph induced by the vertices in the interior of the disk bounded by P. If $\mathbf{tw}(G') > T$, we then G' contains a subdivided wall of height $\kappa + 1$. We set new G' to be the graph induced by the vertices in the strict interior of P in the new subdivided wall and recurse. Otherwise, $\mathbf{tw}(G') \leq T$ and since $\mathbf{tw}(G') \leq t$ G', G' contains a subdivision of a wall of height κ. ◻

Finding an irrelevant edge. We have now built all necessary tools we need in order to find an irrelevant vertex inside a calibrated subdivided wall.

Lemma 11. *There exists a function $f_2 : \mathbb{N} \to \mathbb{N}$ and an algorithm that, with input a rooted Σ-embedded graph (G, S_G) and an integer h, outputs either a tree decomposition of G of width at most $\alpha_\Sigma \cdot f_2(h)$ or an h-irrelevant edge of G where α_Σ is a constant depending on Σ.*

Proof. Set $f_2(h) = (f_1(h)/4)^2 + g(h) \cdot (8 \cdot f_1(h)+1) \cdot 2 \cdot h$ (here g is the function of Lemma 9). We apply first on G the algorithm in Lemma 10 for $\kappa = f_2(h)$. As a result, there exist integers t and T such that we either have that $\mathbf{tw}(G) \leq t$ or we find a subdivided wall $W$of height $f_2(h)$ as a subgraph of G with the property that the graph K induced by the vertices of the interior of the disk (from Lemma 10) has treewidth at most T. Clearly, it is enough to examine the second case.

It follows from Lemma 10 that K contains an (r, q)-railed annulus $(\mathcal{A}, \mathcal{W})$ where $r = g(h) \cdot (8 \cdot f_1(h) + 1) \cdot 2 \cdot h$ and $q = (f_1(h))^2$. Let I be any subseteq $I \subseteq \{1, \ldots, q\}$ with $f_1(k)$ elements. We now apply the FPT algorithm of Lemma 9 for $l = (8 \cdot f_1(h) + 1) \cdot 2 \cdot h$ and obtain some $\theta \in [r]$ such that $\chi_{\theta+1} = \cdots = \chi_{\theta+l}$. We pick an edge e from $E(C_{\theta+2}) \setminus \bigcup_{j \in X_I^{(\theta+2)}} E(P^{(\theta+2,j)})$ and we claim that e is irrelevant. For this, we assume that $\mathcal{P}'$ is a collection of internally disjoint paths such that $G_{\mathcal{P}}$ is ϕ-topologically isomorphic to (H, S_H) such that $r \subseteq \phi$. Our target is to replace $\mathcal{P}'$ by a new one $\mathcal{P}$ such that $v \notin V(\mathcal{P})$. Let Y be the set of the non-internal vertices of the paths in $\mathcal{P}'$. We define $\mathcal{L}_{\mathcal{P}'}$ by removing from the set $\{P[V(P) \setminus Y] \mid P \in \mathcal{P}'\}$ the empty graphs and observe that $\mathcal{L}_{\mathcal{P}'}$ is an ψ-linkage of G for some $\psi \leq |E(H)|$. Let now U be the set of the endpoints of the paths in $\mathcal{L}_{\mathcal{P}'}$ and observe that $|U| \leq 2\psi \cdot \leq 2 \cdot h$. As a consequence of that, we have that there is some θ' such that, if $r' = 8 \cdot f_1(h) + 1$, then the territory of the (r', q)-annulus $\mathcal{A}' = \mathcal{A}[\theta + \theta' + 1, \theta + \theta' + r']$ does not contain vertices from U and thus is $\mathcal{L}_{\mathcal{P}'}$-avoiding. From Lemma 7, $\mathcal{L}_{\mathcal{P}'}$ can be replaced by a equivalent one, we call it $\mathcal{L}^{\text{new}}$ such that $\mathcal{L}^{\text{new}}$ abandons $\Delta^{(y)}$ on $X_I^{(y)}$ where $y = 4 \cdot f_1(h) + 1$. We now consider the Σ-embedded graph $H^\bullet = (G[E(\mathcal{P}') \setminus E(\mathcal{L}_{\mathcal{P}'})] \cup (\bigcup_{L \in \mathcal{L}^{\text{new}}} L))$ and define H^1 (resp. H^2) as the graph obtained if, in $H^\bullet$ we dissolve all vertices that are not in $Y \cup X_I^y$ (resp Y). Clearly, H^1 is topologically isomorphic to H and the graph $J = H^2 \cap \Delta^{(y)}$ is a $\Delta^{(y)}$-embedded graph that can be $\Delta^{(y)}$-excised from H^1. As $|J|$ has at most h edges and the crossing number of $\Delta^{(y)}$ in H^1 is at most $|X_I^y| = f_1(h)$, it follows that $J \in \mathbf{H}_h$. Let now $\rho : V(J) \cap \mathbf{bor}(\Delta^{(y)}) \to X_I^{(i)}$ be the injection obtained by mapping each vertex of $V(J) \cap \mathbf{bor}(\Delta^{(y)})$ to itself. It follows that J is ρ-attached topological minor of $G^{(y)}$ and therefore $\chi_y(J, \rho) = 1$. This implies that $\chi_{\theta+1}(J, \rho) = \chi_y(J, \rho) = 1$, therefore, J is also a ρ'-attached topological minor of $G^{(y)}$ for $\rho' = \lambda_{\theta+1,y} \circ \rho$. This, in turn, implies that there exists a collection $\mathcal{P}^-$ of paths in $G^{(\theta+1)}$ such that the graph obtained by $G[\mathcal{P}^-]$ if we connect by edges all pairs of vertices in $X_I^{(\theta+1)} \cap V(\mathcal{P})$ that are successive in $\Delta^{(\theta+1)}$, is ϕ'-topologically isomorphic to $J_{\Delta^{(\theta+1)}, \Delta^{(\theta+1)} \cap V(J)}$ for some $\phi' \supseteq \rho'$ where J is now drawn inside $\Delta^{(\theta+1)}$. It is now easy to observe that the graph

$$H^\star = G[\mathcal{P}^-] \cup G[\bigcup_{j \in I} \{e \in E(W_j) \mid e \subseteq \Delta^q \setminus \Delta^{\theta+1}\}] \cup G[\{e \in E(H^\bullet) \mid e \subseteq \Sigma \setminus \Delta^q\}]$$

is isomorphic to $G_{\mathcal{P}}$ for some collection $\mathcal{P}$ of internally disjoint paths of G. Moreover $(H^\star, S_H)$ is topologically isomorphic to H for some $\phi'' \supseteq \sigma$ (recall that $S_H \cap \Delta_q = \emptyset$). As in the construction of $H^\star$ all edges of $C_{\theta+1}$ that are used are edges from $\bigcup_{j \in X_I^{(\theta+2)}} P^{(\theta+2,j)}$, we have that $e \notin E(H^\star)$. Thus, e is irrelevant. $\square$

Theorem 1. *For every surface Σ, there exists an FPT algorithm to decide whether a given Σ-embedded graph H, rooted on S_H, is a σ-rooted surface topological minor of a Σ-embedded input graph, rooted on S_G, for some bijection $\sigma : S_H \to S_G$, when parameterized by the size of H.*

Proof. Let (G, S_G) be a Σ-embedded input graph. We apply Lemma 11. If the algorithm returns the tree decomposition, we solve the problem by the standard dynamic programming techniques. (Rooted surface topological minors are expressible in MSOL.) If the algorithm returns, an h-irrelevant vertex v of G, then we run the algorithm for $G \setminus v$. $\square$

The following result is a consequence of Theorem 1 (we take σ to be the void function), Proposition 2, and the fact that set $C_\Sigma(H)$ is finite (Lemma 5 in [15]).

Theorem 2. *For every surface Σ, there exists an* FPT *algorithm to decide whether a given graph H is a contraction of a Σ-embedded input graph, when parameterized by the size of H.*

9 Open problem

We prove that contraction checking is FPT for graphs on surfaces. To complete Table 1 it would be interesting to know the parametrized complexity of induced minor checking for graphs on surfaces.

References

1. Isolde Adler, Frederic Dorn, Fedor V. Fomin, Ignasi Sau, and Dimitrios M. Thilikos. Faster parameterized algorithms for minor containment. In *SWAT*, pages 322–333, 2010.
2. Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.
3. A. E. Brouwer and H. J. Veldman. Contractibility and NP-completeness. *Journal of Graph Theory*, 11(1):71–79, 1987.
4. Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990.
5. Erik D. Demaine, Fedor V. Fomin, Mohammadtaghi Hajiaghayi, and Dimitrios M. Thilikos. Subexponential parameterized algorithms on bounded-genus graphs and H-minor-free graphs. *J. Assoc. Comput. Mach.*, 52(6):866–893, 2005.
6. Reinhard Diestel. *Graph Theory*. Springer-Verlag, Electronic Edition, 2005.
7. R.G. Downey and M. R. Fellows. *Parameterized Complexity*. Springer-Verlag, 1999.
8. David Eppstein. Subgraph isomorphism in planar graphs and related problems. *J. Graph Algorithms Appl.*, 3(3), 1999.
9. Michael R. Fellows, Jan Kratochvíl, Martin Middendorf, and Frank Pfeiffer. The complexity of induced minors and related problems. *Algorithmica*, 13(3):266–282, 1995.
10. J. Flum and M. Grohe. *Parameterized complexity theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, Berlin, 2006.
11. Fedor V. Fomin, Petr A. Golovach, and Dimitrios M. Thilikos. Contraction bidimensionality: The accurate picture. In Amos Fiat and Peter Sanders, editors, *ESA*, volume 5757 of *Lecture Notes in Computer Science*, pages 706–717. Springer, 2009.
12. Fedor V. Fomin, Petr A. Golovach, and Dimitrios M. Thilikos. Contraction obstructions for treewidth. *J. Comb. Theory, Ser. B*, 101(5):302–314, 2011.
13. J. F. Geelen, R. B. Richter, and G. Salazar. Embedding grids in surfaces. *European J. Combin.*, 25(6):785–792, 2004.
14. Martin Grohe, Ken-ichi Kawarabayashi, Dániel Marx, and Paul Wollan. Finding topological subgraphs is fixed-parameter tractable. *To appear in Proceedings of the 43rd ACM Symposium on Theory of Computing(STOC'11)*, 2011.
15. Marcin Kamiński, Daniël Paulusma, and Dimitrios M. Thilikos. Contractions of planar graphs in polynomial time. In Mark de Berg and Ulrich Meyer, editors, *ESA*, volume 6346 of *Lecture Notes in Computer Science*, pages 122–133. Springer, 2010.
16. Marcin Kamiński, Daniël Paulusma, and Dimitrios M. Thilikos. Contractions of planar graphs in polynomial time. *manuscript*, 2011.
17. Jirí Matousek and Robin Thomas. On the complexity of finding iso- and other morphisms for partial k-trees. *Discrete Mathematics*, 108(1-3):343–364, 1992.
18. Frédéric Mazoit. Tree-width of hypergraphs and surface duality. *CoRR*, abs/1006.3167, 2010.
19. B. Mohar and C. Thomassen. *Graphs on Surfaces*. The Johns Hopkins University Press, 2001.
20. B. Mohar and C. Thomassen. *Graphs on surfaces*. John Hopkins University Press, 2001.
21. Bojan Mohar. A linear time algorithm for embedding graphs in an arbitrary surface. *SIAM J. Discrete Math.*, 12(1):6–26, 1999.
22. Rolf Niedermeier. *Invitation to fixed-parameter algorithms*, volume 31 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford University Press, Oxford, 2006.
23. Neil Robertson and Paul D. Seymour. Graph Minors. XXII. Irrelevant vertices in linkage problems. *Journal of Combinatorial Theory, Series B.* (to appear).
24. Neil Robertson and Paul D. Seymour. Graph Minors XIII. The disjoint paths problem. *J. Comb. Theory, Ser. B*, 63(1):65–110, 1995.
25. Neil Robertson and Paul D. Seymour. Graph Minors. XXI. Graphs with unique linkages. *Journal of Combinatorial Theory, Series B*, 99(3):583–616, 2009.