Excluding Graphs as Immersions in Surface Embedded Graphs

Archontia C. Giannopoulou[†¶ ‖], Marcin Kamiński[‡],
and Dimitrios M. Thilikos[§∗∗]

Abstract

We prove a structural characterization of graphs that forbid a fixed
graph H as an immersion and can be embedded in a surface of Eüler genus
γ. In particular, we prove that a graph G that excludes some connected
graph H as an immersion and is embedded in a surface of Eüler genus
γ has either "small" treewidth (bounded by a function of H and γ) or
"small" edge connectivity (bounded by the maximum degree of H). Using
the same techniques we also prove an excluded grid theorem on bounded
genus graphs for the immersion relation.

Keywords: Surface Embeddable Graphs, Immersion Relation, Treewidth, Edge
Connectivity.

1   Introduction

A graph H is an immersion of a graph G if it can be obtained from G by
removing vertices or edges, and splitting off adjacent pairs of edges. The class of
all graphs was proved to be well-quasi-ordered under the the immersion relation
by Robertson and Seymour in the last paper of their Graph Minors series [20].
Certainly, this work was mostly dedicated to minors and not immersions and has
been the source of many theorems regarding the structure of graphs excluding
some graph H as a minor. Moreover, the minor relation has been extensively

studied the past two decades and many structural results have been proven for minors with interesting algorithmic consequences (see, for example, [4,14,17–19, 21]). However, structural results for immersions started appearing only recently. In 2011, DeVos et al. proved that if the minimum degree of a graph G is $200t$, then G contains the complete graph on t vertices as an immersion [5]. In [7] Ferrara et al., provided a lower bound (depending on graph H) on the minimum degree of a graph G that ensures that H is contained in G as an immersion. Furthermore, Wollan recently proved a structural theorem for graphs excluding complete graphs as immersions as well as a sufficient condition such that any graph which satisfies the condition admits a wall as an immersion [22]. The result in [22] can be seen as an immersion counterpart of the grid exclusion theorem [17], stated for walls instead of grids and using an alternative graph parameter instead of treewidth.

In terms of graph colorings, Abu-Khzam and Langston in [1] provided evidence supporting the immersion ordering analog of Hadwiger's Conjecture, that is, the conjecture stating that if the chromatic number of a graph G is at least t, then G contains the complete graph on t vertices as an immersion, and proved it for $t \leq 4$. For $t = 5, 6, 7$, see [6, 15]. For algorithmic results on immersions, see [2, 10, 12, 13].

In this paper, we prove structural results for the immersion relation on graphs embeddable on a fixed surface. In particular, we show that if G is a graph that is embeddable on a surface of Eüler genus γ and H is a connected graph then one of the following is true: either G has bounded treewidth (by a function that depends only on γ and H), or its edge connectivity is bounded by the maximum degree of H, or it contains H as a (strong) immersion. Furthermore, we refine our results to obtain a counterpart of the grid exclusion theorem for immersions. In particular, we prove (Theorem 3) that there exists a function $f : \mathbb{N} \to \mathbb{N}$ such that if G is a 4-edge-connected graph embedded on a surface of Eüler genus γ and the treewidth of G is at least $f(\gamma) \cdot k$, then G contains the $k \times k$-grid as an immersion. Notice that the edge connectivity requirement is necessary here as big treewidth alone is not enough to ensure the existence of a graph with a vertex of degree 4 as an immersion. Although a wall of height at least h has treewidth at least h, it does not contain the complete graph on t vertices as an immersion, for any $t \geq 5$. Finally, our results imply that when restricted to graphs of sufficiently big treewidth embeddable on a fixed surface, large edge connectivity forces the existence of a large clique as an immersion.

Our result reveals several aspects of the behavior of the immersion relation on surface embeddable graphs. The proofs exploit variants of the grid exclusion theorem for surfaces proved in [8] and [11] and the results of Biedl and Kaufmann [3] on optimal orthogonal drawings of graphs.

The paper is organized as follows. In Section 2 we give some basic definitions and preliminaries. In Section 3 we give a series main combinatorial results. Based on the results of Section 3, we prove the main theorem and we derive its corollaries in Section 4.

2

2   Preliminaries

For every positive integer n, let $[n]$ denote the set $\{1, 2, \ldots, n\}$. A *graph* G is a pair (V, E) where V is a finite set, called the *vertex set* and denoted by $V(G)$, and E is a set of 2-subsets of V, called the *edge set* and denoted by $E(G)$. If we allow E to be a multiset then G is called a multigraph. Let G be a graph. For a vertex v, we denote by $N_G(v)$ its *(open) neighborhood*, that is, the set of vertices which are adjacent to v, and by $E_G(v)$ the set of edges containing v. Notice that if G is a multigraph $|N_G(v)| \leq |E_G(v)|$. The degree of a vertex v is $\deg_G(v) = |E_G(v)|$. We denote by $\Delta(G)$ the maximum degree over all vertices of G.

If $U \subseteq V(G)$ (respectively $u \in V(G)$ or $E \subseteq E(G)$ or $e \in E(G)$) then $G - U$ (respectively $G - u$ or $G - E$ or $G - e$) is the graph obtained from G by the removal of vertices of U (respectively of vertex u or edges of E or of the edge e). We say that a graph H is a subgraph of a graph G, denoted by $H \subseteq G$, if H can be obtained from G after deleting edges and vertices.

We say that a graph H is an *immersion* of a graph G (or H is *immersed* in G), $H \leq_{\mathrm{im}} G$, if there is an injective mapping $f : V(H) \to V(G)$ such that, for every edge $\{u, v\}$ of H, there is a path from $f(u)$ to $f(v)$ in G and for any two distinct edges of H the corresponding paths in G are *edge-disjoint*, that is, they do not share common edges. The function f is called a *model of H in G*.

Let P be a path and $v, u \in V(P)$. We denote by $P[v, u]$ the subpath of P with endvertices v and u. Given two paths P_1 and P_2 who share a common endpoint v, we say that they are *well-arranged* if their common vertices appear in the same order in both paths.

A *tree decomposition* of a graph G is a pair (T, B), where T is a tree and B is a function that maps every vertex $v \in V(T)$ to a subset B_v of $V(G)$ such that:

(i) $\bigcup_{v \in V(T)} B_v = V(G)$,

(ii) for every edge e of G there exists a vertex t in T such that $e \subseteq B_t$, and

(iii) for every $v \in V(G)$, if $r, s \in V(T)$ and $v \in B_r \cap B_s$, then for every vertex t on the unique path between r and s in T, $v \in B_t$.

The width of a tree decomposition (T, B) is $\mathrm{width}(T, B) := \max\{|B_v| - 1 \mid v \in V(T)\}$ and the treewidth of a graph G is the minimum over the $\mathrm{width}(T, B)$, where (T, B) is a tree decomposition of G.

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 a 2-cell embedding of G in Σ. To simplify notations, 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. 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 $H \subseteq G$ in Σ. Recall that $\Delta \subseteq \Sigma$ is an open (respectively closed) disc if it is homeomorphic to $\{(x, y) : x^2 + y^2 < 1\}$ (respectively $\{(x, y) : x^2 + y^2 \le 1\}$). The *Eüler genus* of a non-orientable surface Σ is equal to the non-orientable genus $\tilde{g}(\Sigma)$ (or the crosscap number). The *Eüler genus* of an orientable surface Σ is $2g(\Sigma)$, where $g(\Sigma)$ is the orientable genus of Σ. We refer to the book of Mohar and Thomassen [16] for more details on graphs embeddings. The *Eüler genus* of a graph G (denoted by $\mathbf{eg}(G)$) is the minimum integer γ such that G can be embedded on a surface of the Eüler genus γ.

Walls. Let k and r be positive integers where $k, r \ge 2$. The $(k \times r)$-*grid* $\Gamma_{k,r}$ is the Cartesian product of two paths of lengths $k - 1$ and $r - 1$ respectively. A *wall of height k*, $k \ge 1$, is the graph obtained from a $((k + 1) \times (2 \cdot k + 2))$-grid with vertices (x, y), $x \in \{1, \ldots, 2 \cdot k + 4\}$, $y \in \{1, \ldots, k + 1\}$, 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. We denote such a wall by W_k. The *corners* of the wall W_k are the vertices $c_1 = (1, 1)$, $c_2 = (2 \cdot k + 1, 1)$, $c_3 = (2 \cdot k + 1 + (k + 1 \mod 2), k + 1)$ and $c_4 = (1 + (k + 1 \mod 2), k + 1)$. (The square vertices in Figure 1.)

A *subdivided wall W of height k* is a wall obtained from W_k after replacing some of its edges by paths without common internal vertices. We call the resulting graph W a *subdivision* of W_k and the new vertices *subdivision vertices*. The non-subdivision vertices are called *original*. The *perimeter P* of a subdivided wall (grid) is the cycle defined by its boundary.

Let W be a subdivided wall in a graph G and K' be the connected component of $G \setminus P$ that contains $W \setminus P$. The *compass K* of W in G is the graph $G[V(K') \cup V(P)]$. Observe that W is a subgraph of K and K is connected.

The *layers* of a subdivided wall W of height k are recursively defined as follows. The first layer of W, denoted by L_1, is its perimeter. For $i = 2, \cdots, \lceil \frac{k}{2} \rceil$, the i-th layer of W, denoted by L_i, is the $(i - 1)$-th layer of the subwall W' obtained from W after removing from W its perimeter and (recursively) all occurring vertices of degree 1 (see Figure 1).
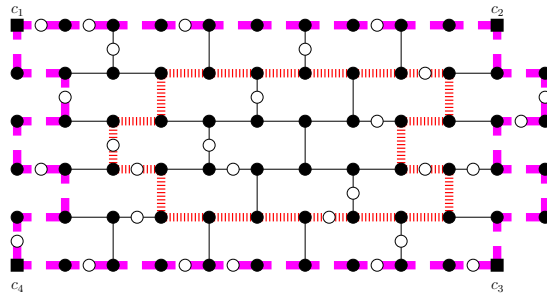

Figure 1: The first (magenta-dashed) and second (red-dotted) layers of a wall of height 5

We denote by A_i the annulus defined by the cycles L_i and L_{i+1}, that is, by i-th and $(i+1)$-th layer, $i \in [[\frac{k}{2}] - 1]$. Given an annulus A defined by two cycles C_1 and C_2, we denote by A° the interior of A, that is, $A \setminus (C_1 \cup C_2)$.

A subdivided wall of height k is called *tight* if

1. the closed disk defined by the innermost ($\lceil \frac{k}{2} \rceil$-th) layer of W is edge-maximal (for reasons of uniformity we will denote this disk by $A_{\lceil \frac{k}{2} \rceil}$),

2. for every $i \in [[\frac{k}{2}] - 1]$ the annulus A_i is edge-maximal under the condition that A_{i+1} is edge-maximal.

If W is a subdivided wall of height k, we call *brick* of W any facial cycle whose non-subdivided counterpart in W_h has length 6. We say that two bricks are *neighbors* if their intersection contains an edge.

Let W_k be a wall. We denote by $P_j^{(h)}$ the shortest path connecting vertices $(1, j)$ and $(2 \cdot k + 2, j)$ and call these paths the *horizontal paths of W_k*. Note that these paths are vertex-disjoint. We call the paths $P_{k+1}^{(h)}$ and $P_1^{(h)}$ the *southern path of W_k* and *northern part of W_k* respectively.

Similarly, we denote by $P_i^{(v)}$ the shortest path connecting vertices $(i, 1)$ and $(i, k+1)$ with the assumption that for, $i < 2 \cdot k + 2$, $P_i^{(v)}$ contains only vertices (x, y) with $x = i, i + 1$. Notice that there exists a unique subfamily $\mathcal{P}_v$ of $\{P_i^{(v)} \mid i < 2 \cdot k + 2\}$ of $k + 1$ vertical paths with one endpoint in the southern path of W_k and one in the northern path of W_k. We call these paths *vertical paths of W_k* and denote them by $P_i^{[v]}$, $i \in [k]$, where $P_1^{(v)} = P_1^{[v]}$ and $P_{2 \cdot k + 1}^{(v)} = P_{k+1}^{[v]}$. (See Figure 2.)
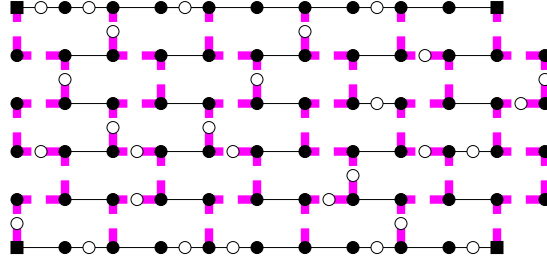

Figure 2: The vertical paths of a wall of height 5

The paths $P_1^{[v]}$ and $P_{k+1}^{[v]}$ are called the *western part of W_k* and the *eastern part of W_k* respectively. Note that the perimeter of the wall can alternatively be defined as the cycle $P_1^h \cup P_{k+1}^h \cup P_1^{[v]} \cup P_{k+1}^{[v]}$.

Notice now that each vertex $u \in V(W_k) \setminus V(P)$, is contained in exactly one vertical path, denoted by $P_u^{(v)}$, and in exactly one horizontal path, denoted by $P_u^{(h)}$, of W_k. If W is a subdivision of W_k, we will use the same notation for the paths obtained by the subdivisions of the corresponding paths of W_k, with

further assumption that u is an original vertex of W.

Given a wall W and a layer L of W, different from the perimeter of W. Let W' be the subwall of W with perimeter L. W' is also called the *subwall of W defined by L*. We call the following vertices, *important* vertices of L; the original vertices of W that belong to L and have degree 2 in the underlying non-subdivided wall of W' but are not the corners of W' (where we assume that W' shares the original vertices of W). (See Figure 3)
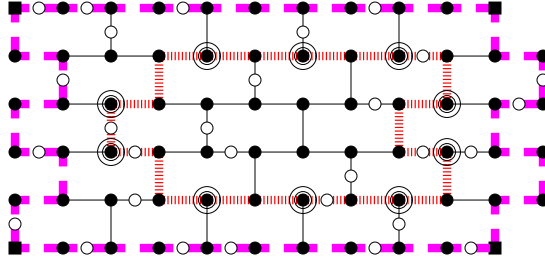

Figure 3: The important vertices the second layer of a wall of height 5

Observation 1. *A layer L of a wall W that is different from its perimeter and defines a subwall W' of W of height k contains exactly $4k-2$ important vertices.*

From Lemma 6 in [8] and Lemma 3 in [11] we obtain the following.

Lemma 1. *Let G be a graph embedded in a surface of Eüler genus γ. If $\mathbf{tw}(G) \geq 48 \cdot (\gamma + 1)^{\frac{3}{2}} \cdot (k + 5)$, G contains as a subgraph a subdivided wall of height k, whose compass in G is embedded in a closed disk Δ.*

Confluent paths. Let G be a graph embedded in some surface Σ and let $x \in V(G)$. We define a *disk around x* as any open disk Δ_x with the property that each point in $\Delta_x \cap G$ is either x or belongs to the edges incident to x. Let P_1 and P_2 be two edge-disjoint paths in G. We say that P_1 and P_2 are *confluent* if for every $x \in V(P_1) \cap V(P_2)$, that is not an endpoint of P_1 or P_2, and for every disk Δ_x around x, one of the connected components of the set $\Delta_x \setminus P_1$ does not contain any point of P_2. We also say that a collection of paths is *confluent* if the paths in it are pairwise confluent.

Moreover, given two edge-disjoint paths P_1 and P_2 in G we say that a vertex $x \in V(P_1) \cap V(P_2)$ that is not an endpoint of P_1 or P_2 is an *overlapping vertex of P_1 and P_2* if there exists a Δ_x around x such that both connected components of $\Delta_x \setminus P_1$ contain points of P_2. (See, Figure 4.) For a family of paths $\mathcal{P}$, a vertex v of a path $P \in \mathcal{P}$ is called an *overlapping vertex of $\mathcal{P}$* if there exists a path $P' \in \mathcal{P}$ such that v is an overlapping vertex of P and P'.
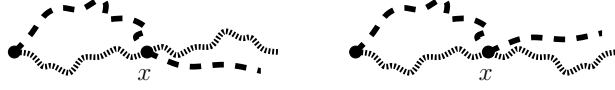
Figure 4: The vertex x is an overlapping vertex of the two paths on the left (dashed and dotted), while it is not an overlapping vertex of the paths on the right.

Orthogonal drawings. An *orthogonal drawing* of a graph G in a grid Γ is a mapping which maps vertices $v \in V(G)$ to subgrids $\Gamma(v)$ (called *boxes*) such that for every $u_1, u_2 \in V(G)$ with $u_1 \neq u_2$, $\Gamma(u_1) \cap \Gamma(u_2) = \emptyset$, and edges $\{u_1, u_2\} \in E(G)$ to (u_1', u_2')-paths whose internal vertices belong to $\Gamma - \bigcup_{v \in V(G)} \Gamma(v)$, their endpoints u_i' (called *joining vertices of $\Gamma(u_i)$*) belong to the perimeter of $\Gamma(u_i)$, $i \in [2]$, and for every two disjoint edges $e_i \in E(G)$, $i \in [2]$, the corresponding paths are edge-disjoint.

We need the following result.

Lemma 2 ([3]). *If G is a simple graph then it admits an orthogonal drawing in an $(\frac{m+n}{2} \times \frac{m+n}{2})$-grid. Furthermore, the box size of each vertex v is $\frac{\deg(v)+1}{2} \times \frac{\deg(v)+1}{2}$.*

3 Preliminary Combinatorial Lemmata

Before proving the main result of this section we first state the following lemma which we will need later on.

Lemma 3 ([9]). *Let r be a positive integer. If G is a graph embedded in a surface Σ, $v, v_1, v_2, \ldots, v_r \in V(G)$, and $\mathcal{P}$ is a collection of r edge-disjoint paths from v to $v_1, v_2, \ldots, v_r$ in G, then G contains a confluent collection $\mathcal{P}'$ of r edge-dsjoint paths from v to $v_1, v_2, \ldots, v_r$ such that $E(\bigcup_{P \in \mathcal{P}'} P) \subseteq E(\bigcup_{P \in \mathcal{P}} P)$.*

Detachment tree of $\mathcal{P}$ in u. Let G be a graph embedded in a closed disk Δ, $v, v_1, v_2, \ldots, v_k$ be distinct vertices of G, and $\mathcal{P} = \{P_i \mid i \in [k]\}$ be a family of k confluent edge-disjoint paths such that P_i is a path from v to v_i, $i \in [k]$. Let also $u \in V(G) \setminus \{v, v_i \mid i \in [k]\}$ be an internal vertex of at least two paths in $\mathcal{P}$. Let $\mathcal{P}_u = \{P_{i_1}, P_{i_2}, \ldots, P_{i_r}\}$ denote the family of paths in $\mathcal{P}$ that contain u and Δ_u be a disk around u. Given any edge e with $u \in e$ we denote by u_e its common point with the boundary of Δ_u. Moreover, we denote by $e_{i_r}^1$ and $e_{i_r}^2$ the edges of P_{i_j} incident to u, $j \in [r]$.

We construct a tree T_u in the following way and call it *detachment tree of $\mathcal{P}$ in u*. Consider the outerplanar graph obtained from the boundary of Δ_u by adding the edges $\{u_{e_{i_j}^1}, u_{e_{i_j}^2}\}$, $j \in [r]$. We subdivide the edges $\{u_{e_{i_j}^1}, u_{e_{i_j}^2}\}$, $j \in [r]$, resulting to a planar graph. For every bounded face f of the graph, let $V(f)$ denote the set of vertices that belong to f. We add a vertex v_f in its interior and

we make it adjacent to the vertices of $(V(f) \cap \{u_e \mid e \in u\}) \setminus \{u^1_{e_{i_j}}, u^2_{e_{i_j}} \mid j \in [r]\}$. Finally we remove the edges that lie in the boundary of Δ_u. We call this tree T_u. Notice that for every e with $u \in e$, the vertex u_e is a leaf of T_u. (See Figure 5.)

We replace u by T_u in the following way. We first subdivide every edge $e \in G$ incident to u, and denote by u_e the vertex added after the subdivision of the edge e. We denote by G_s the resulting graph. Consider now the graph $G^r = (G^s \setminus u) \cup T_u$ (where, without loss of generality, we assume that $V(G \setminus u) \cap V(T_u) = \{u_e \mid u \in e\}$). The graph G^r is called *the graph obtained from G by replacing u with T_u*.
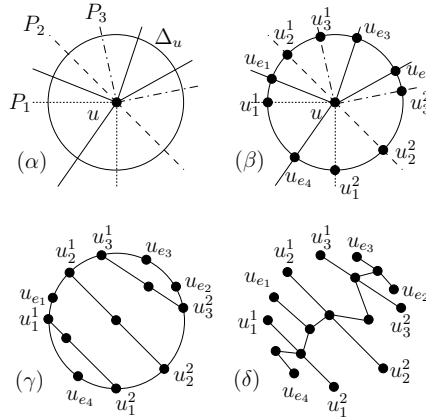

Figure 5: Example of the construction of a detachment tree.

Observation 2. *Let k, h be positive integers and G be a multigraph containing as a subgraph a subdivided wall W of height h, whose compass C is embedded in a closed disk Δ. Furthermore, let $v, v_i, i \in [k]$, be vertices of W such that there exists a confluent family $\mathcal{P}$ of k edge-disjoint paths from v to the vertices $v_i, i \in [k]$. Finally, let $u \in V(C) \setminus \{v, v_i \mid i \in [k]\}$ belonging to more than one paths of $\mathcal{P}$. The graph G^r obtained from G by replacing u with T_u contains as a subgraph a subdivided wall W' of height h, whose compass is embedded in Δ and there exists a family $\mathcal{P}'$ of k confluent edge-disjoint paths from v to $v_i, i \in [k]$, in W' whose paths avoid u.*

Proof. Notice first that it is enough to prove the observation for the case where $u \in V(W)$. Let e_1, e_2 (and possibly e_3) be the edges incident to u that also belong to W. Notice now that the vertices u_{e_1}, u_{e_2} (and u_{e_3}) are leaves of T_u. Thus, from a folklore result, there exists a vertex $u' \in V(T_u)$ such that there exist 2 (or 3) internally vertex-disjoint paths from u' to u_{e_1} and u_{e_2} (and possibly u_{e_3}). □

We now state the following auxiliary definitions. Let G be a multigraph that

contains a wall of height k whose compass is embedded in a closed disk. Let $v \in A_{\lceil \frac{k}{2} \rceil}$, that is, let v be a vertex contained in the closed disk defined by the innermost layer of W, and let P be a path from v to the perimeter of W. For each layer j of the wall, $2 \le j \le \lceil \frac{k}{2} \rceil$, we denote by x_P^j the first vertex of P (starting from v) that also belongs to L_j and we call it *incoming vertex of P in L_j*.

We denote by P^j the maximal subpath of P that contains v and is entirely contained in the wall defined by L_j. Moreover, we denote by y_P^j its endpoint in L_j and call it *outgoing vertex of P in L_j*. Notice that x_P^j and y_P^j are not necessarily distinct vertices.

Lemma 4. *Let λ and k be positive integers. Let G be a graph and W be a tight subdivided wall of G of height k, whose compass is embedded in a closed disk Δ. Let also v be a vertex such that $v \in A_{\lceil \frac{k}{2} \rceil}$. If there exist λ vertex-disjoint paths P_i, $i \in [\lambda]$, from v to vertices of the perimeter then there is a brick B of W with $B \cap A_{j-1}^{o} \ne \emptyset$ that contains both $y_{P_i}^{j}$ and $x_{P_i}^{j-1}$.*

Proof. Assume the contrary. Then it is easy to see that we can construct an annulus A_j' such that $A_j \subsetneq A_j'$ and $|E(A_j)| < |E(A_j')|$, a contradiction to the tightness of the wall. (See Figure 6.) $\qquad\square$
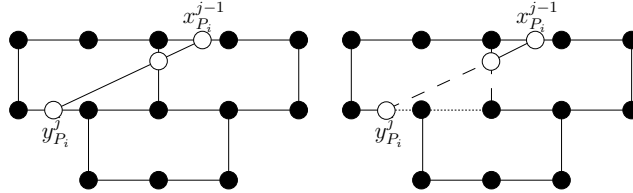

Figure 6: We replace the dotted line of the wall by the dashed line.

Lemma 5. *Let k be a positive integer and G be a multigraph that contains as a subgraph a subdivided wall W of height at least $4 \cdot k^2 + 1$, whose compass K is embedded in a closed disk Δ. Let also V be a set of k vertices lying in the perimeter P of W, whose mutual distance in the underlying non-subdivided wall is at least 2. If there exist a vertex $v \in A_{2 \cdot k^2 + 1}$ and k internally vertex-disjoint paths from v to vertices of P, then there exist k internally vertex-disjoint paths from v to the vertices of V in K.*

Proof. Assume first, without loss of generality, that the wall W is tight. Let then $P_1, P_2, \ldots, P_k$ be the paths from v to P and let $[P_1, P_2, \ldots, P_k, P_1]$ be the clockwise cyclic ordering according to which they appear in W. Our objective is to reroute the paths $P_i, i \in [k]$, so that they end up to the vertices of V. To do so our first step is to identify a layer of the wall for which there exist two consecutive paths whose incoming vertices on the layer are "sufficiently far apart".

9

Let $j_0 = k^2 + 1$. Consider the layer L_{j_0} and for every $i \in [k]$ let T_i denote the path of L_{j_0} starting from $x_i^{j_0}$ and ending in $x_{i+1}^{j_0}$ (considered clockwise), that is, the path of L_{j_0} starting from the incoming vertex of P_i in L_{j_0} and ending to the incoming vertex of P_{i+1} in L_{j_0}, where in the case $i = k$ we abuse notation and assume that $x_{k+1}^{j_0} = x_1^{j_0}$ (see Figure 7). Let also $i_0 \in [k]$ be the index such that the path T_{i_0} contains the maximum number of important vertices amongst the T_i's. Without loss of generality we may assume that $i_0 = \lceil \frac{k}{2} \rceil$. From Observation 1, as L_{j_0} defines a subwall of W of height $2k^2 + 1$, L_{j_0} contains exactly $8k^2 + 2$ important vertices. Thus, at least $7k$ important vertices are internally contained in T_{i_0}. This concludes the first step of the proof.
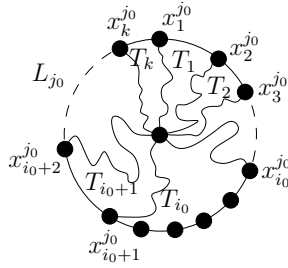

Figure 7: The T_i's, $i \in [k]$

Let now $j_1 = k + 1$. At the next step, using the part of the wall that is contained in $A[L_{j_0}, L_{j_1}]$, that is, in the annulus between the j_0-th and the j_1-th layer of the wall, we find k internally vertex-disjoint paths from the incoming vertices of the paths in L_{j_0} to k consecutive important vertices of the $k + 1$-th layer of the wall. These are the paths that will allow us to reroute the original paths.

Continuing the proof, let $u_1, u_2, \ldots, u_k$ be a set of successive important vertices appearing clockwise in T_{i_0} such that the paths $T_{i_0}[x_{i_0}^{j_0}, u_1]$ and $T_{i_0}[u_k, x_{i_0+1}^{j_0}]$ internally contain at least $3k$ important vertices. Notice that, without loss of generality, we may assume that the vertices u_i, $i \in [k]$, belong to the northern part of W'. Recall here that each original vertex w of $W' \subseteq W \setminus P$ is contained in exactly one vertical path $P_w^{(v)}$ of W. For every $i \in [k]$ we assign the path R_i to the vertex u_i in the following way. Let R_i be the maximal subpath of $P_{u_i}^{(v)}$ whose endpoints are u_i and the important vertex of L_{j_1} that also belongs to $P_{u_i}^{(v)}$, which from now on we will denote by u_i^f. Note here that the paths R_i, $i \in [k]$, are vertex-disjoint and do not contain any of the vertices belonging to the interior of the disk defined by L_{j_0} in the compass of W (See, for example, Figure 8).

Notice now that T_{i_0}, and thus L_{j_0}, contains a path F_1 from $x_{i_0}^{j_0}$ to u_{i_0} and a path F_2 from u_{i_0+1} to $x_{i_0+1}^{j_0}$ that are vertex-disjoint and do not contain vertices of any path other than P_{i_0} and P_{i_0+1}. Consider now the $\lceil \frac{k-2}{2} \rceil$ consecutive layers of W preceding L_{j_0}, that is, the layers $L_j' = L_{j_0-j}$, $j \in [\lceil \frac{k-2}{2} \rceil]$. For every
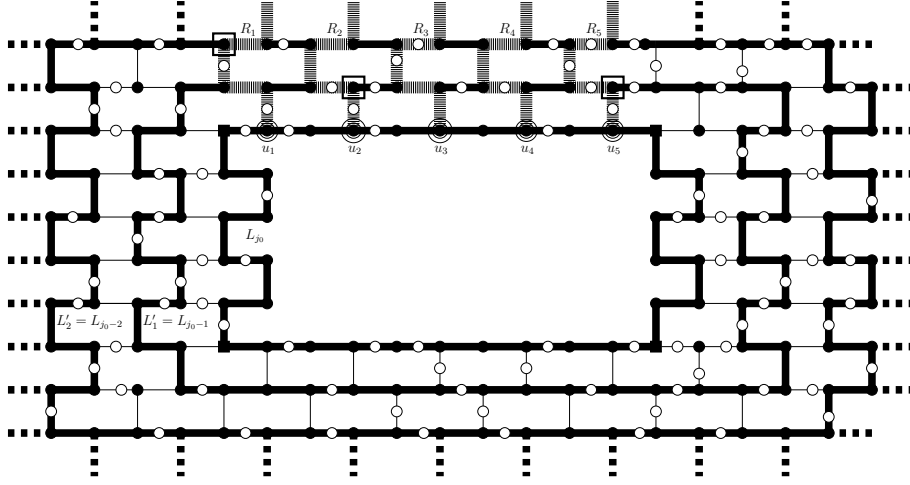
Figure 8: The important vertices of L_{j_0}, the layers L'_1 and L'_2, and the paths R_i

$j \in [\lceil \frac{k-2}{2} \rceil]$ let $u^j_{i_0-j}$ be the first time the path R_{i_0-j} meets L'_j starting from u_{i_0-j} and $u^j_{i_0+1+j}$ be the first time the path R_{i_0+1+j} meets L'_j starting from u_{i_0+1+j}. (See, for example, the vertices inside the squares in Figure 8.)

We need to prove the following.

Claim: For every $j \in [\lceil \frac{k-2}{2} \rceil]$, there exist two vertex-disjoint paths F^1_j and F^2_j between the pairs of vertices $(x^{j_0-j}_{i_0-j}, u^j_{i_0-j})$ and $(x^{j_0-j}_{i_0+1+j}, u^j_{i_0+1+j})$ that do not intersect the paths $\{R_l \mid i_0 - j < l < i_0 + 1 + j\}$.

Proof of Claim: Indeed, this holds by inductively applying the combination of Lemma 4 with the assertion that for every $j \le 2 \cdot k^2 + 1$ and every p, q with $1 < p < q < k$, the outgoing vertices of P_{p-1} and P_{q+1} and the incoming vertices of P_p and P_q in the layer L_j, y^j_{p-1}, y^j_{q+1}, x^j_p, and x^j_q respectively appear in L_j respecting the clockwise order

$$[y^j_{p-1}, x^j_p, x^j_q, y^j_{q+1}]$$

in the tight wall W. This completes the proof of the claim.   ◇

We now construct the following paths. First, let

$$Q_{i_0} = F_1 \cup R_{i_0}$$

and

$$Q_{i_0+1} = F_2 \cup R_{i_0+1},$$

11

that is, Q_{i_0} is the union of the paths F_1 and R_{i_0}, and Q_{i_0+1} is the union of the paths F_2 and R_{i_0+1}. Then, for every $j \in [\lceil \frac{k-2}{2} \rceil]$, let

$$Q_{i_0-j} = P_{i_0-j}[x_{i_0-j}^{j_0}, x_{i_0-j}^{j_0-j}] \cup F_j^1 \cup R_{i_0-j}[u_{i_0-j}^j, u_{i_0-j}^f],$$

that is, Q_{j_0-j} is the union of the following paths; (a) the subpath of P_{i_0-j} between its incoming vertex in the j_0-th layer and its incoming vertex in the (j_0-j)-th layer, (b) the path F_j^1 defined in the claim above, and (c) the subpath of R_{i_0-j} between the vertices $u_{i_0-j}^j$ and $u_{i_0-j}^f$.

Finally, for every $j \in [\lceil \frac{k-2}{2} \rceil]$ ($j \in [\lceil \frac{k-2}{2} \rceil - 1]$, if k is odd) let

$$Q_{i_0+1+j} = P_{i_0+1+j}[x_{i_0+1+j}^{j_0}, x_{i_0+1+j}^{j_0-j}] \cup F_j^2 \cup R_{i_0+1+j}[u_{i_0+1+j}^j, u_{i_0+1+j}^f],$$

that is, Q_{i_0+1+j} is the union of the following three paths; (a) the subpath of P_{i_0+1+j} between its incoming vertex in the j_0-th layer and its incoming vertex in the (j_0-j)-th layer, (b) the path F_j^2 defined in the claim above, and (c) the subpath of R_{i_0+1+j} between the vertices $u_{i_0+1+j}^j$ and $u_{i_0+1+j}^f$.

From the claim above and Lemma 4 we get that the above paths are vertex-disjoint. This concludes the second step of the proof.

We claim now that we may reroute the paths P_i, $i \in [k]$, in such a way that they end up to the vertices u_i^f, $i \in [k]$. Indeed, let $P_i' = P_i[v, x_i^{j_0}] \cup Q_i$, $i \in [k]$. From their construction these paths are vertex-disjoint and end up to the vertices u_i^f, $i \in [k]$. (For a rough estimation of the position of the paths in the wall see Figure 9.)
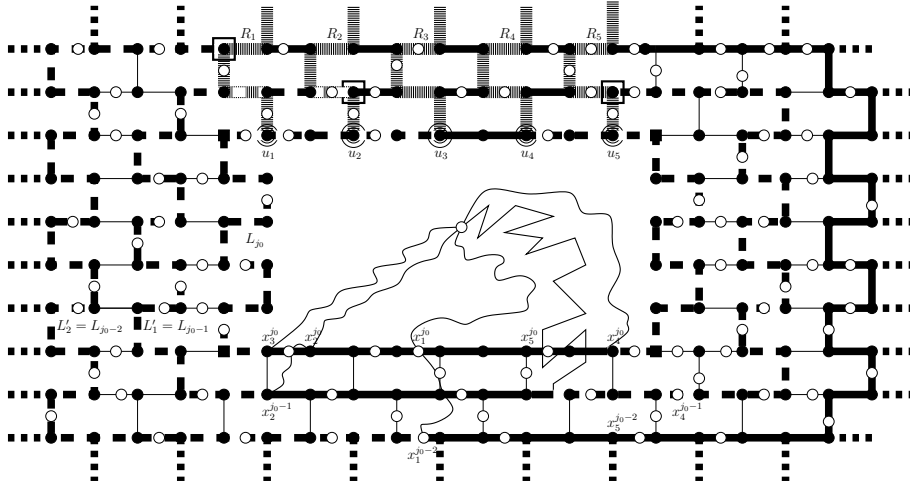

Figure 9: Part of the rerouted paths

Concluding the proof, as the mutual distance of the vertices of V in the underlying non-subdivided wall is at least 2, it is easy to notice that in the

12

annulus defined by L_{j_1} and L_1 there exist k vertex-disjoint paths from the vertices u_i^f, $i \in [k]$, to the vertices of V. $\qquad\square$

We now prove the main result of this section.

Lemma 6. *Let k be a positive integer and G be a k-edge-connected multigraph embedded in a surface of Eüler genus γ that contains a subdivided wall W of height at least $4 \cdot k^2 + 1$ as a subgraph, whose compass C is embedded in a closed disk Δ. Let also S be a set of vertices in the perimeter of W whose mutual distance in the underlying non-subdivided wall is at least 2. If $|S| \leq k$ then there exist a vertex v in W and $|S|$ edge-disjoint paths from v to the vertices of S.*

Proof. Let $v \in A_{2k^2+1}$ and $u \in L_1$ be vertices belonging to the closed disk defined by the layer L_{2k^2+1} and to the perimeter of the wall respectively. As G is k-edge-connected there exist k edge-disjoint paths $P_1, P_2, \ldots, P_k$ connecting v and u. By Lemma 3, we may assume that the paths are confluent. Let $\mathcal{P}' = \{P_i' \mid i \in [k]\}$ be the family of paths $P_i' = P_i[v, x_i^1]$, $i \in [k]$, that is, let $\mathcal{P}'$ be the family of paths consisting of he subpaths of P_i, $i \in [k]$, between v and the first vertex on which they meet the perimeter of W.

Let V be the set of vertices in $V(C) \setminus (V(L_1) \cup \{v\})$ that are contained in more than one path in $\mathcal{P}'$. We obtain the graph $\hat{G}$ by replacing every vertex $z \in V$ with the detachment tree of $\mathcal{P}'$ in z. From Observation 2, $\hat{G}$ contains a wall $\hat{W}$ of height $4k^2 + 1$ whose compass is embedded in Δ. Notice also that, as no changes have occurred in the perimeter of W, W and $\hat{W}$ share the same perimeter. Furthermore, $\hat{W}$ contains k internally vertex-disjoint paths from v to the perimeter of $\hat{W}$. Thus, from Lemma 5, $\hat{W}$ contains k vertex-disjoint paths from v to S. It is now easy to see, by contracting each one of the trees T_z, $z \in V(C) \setminus (V(L_1) \cup \{v\})$, to a single vertex that W contains k edge-disjoint paths from v to S. $\qquad\square$

4 Main Theorem

By combining Lemmata 6, 1 and 2 we obtain the following.

Theorem 1. *There exists a computable function $f : \mathbb{N} \to \mathbb{N}$ such that for every multigraph G of Eüler genus γ and every connected graph H one of the following holds:*

1. $\mathbf{tw}(G) \leq f(\gamma) \cdot \lambda \cdot k$, *where $\lambda = \Delta(H)$ and $k = \mathbf{m}(H)$*

2. *G is not λ-edge-connected,*

3. *$H \leq_{im} G$.*

Proof. Let
$$ f(\gamma, \lambda, k) = 48 \cdot (\gamma + 1)^{\frac{3}{2}} \cdot \left(\frac{4\,(4\lambda + 1)\,k}{2} + 5 \right), $$

13

and assume that $\mathbf{tw}(G) \geq f(\gamma, \lambda, k)$ and G is λ-edge-connected. From Lemma 1, we obtain that G contains as a subgraph a subdivided wall W of height $2(2\lambda+1)k$ whose compass is embedded in a closed disk.

In what follows we will construct a model of H into the wall. From Lemma 2, H admits an a orthogonal drawing ψ in an

$$\left(\frac{\mathbf{m}(H) + \mathbf{n}(H)}{2} \times \frac{\mathbf{m}(H) + \mathbf{n}(H)}{2} \right) \text{-grid,}$$

where the box size of each vertex $v \in V(H)$ is $\frac{\deg(v)+1}{2} \times \frac{\deg(v)+1}{2}$.

Notice now that ψ can be scaled to an orthogonal drawing ϕ to the grid Γ of size

$$\left(\frac{2\,(4\lambda + 1)\,(\mathbf{m}(H) + \mathbf{n}(H))}{2} + 1 \right) \times 2 \left(\frac{2\,(4\lambda + 1)\,(\mathbf{m}(H) + \mathbf{n}(H)) + 2}{2} + 1 \right),$$

where the box size of each vertex is $(4(\deg(v))^2 + 2) \times 2(4(\deg(v))^2 + 2)$, the joining vertices of each box have mutual distance at least 2 in the perimeter of the box and no joining vertex is a corner of the box.

Moreover, for every vertex u, $u \in \mathbf{Im}(\phi) \setminus \cup_{v \in V(H)} \Gamma(v)$ of degree 4, that is, for every vertex in the image of ϕ that is the intersection of two paths, there is a box in the grid of size $(4 \deg(u)^2 + 2) \times 2(4 \deg(u)^2 + 2)$, denoted by $Q(u)$, containing only this vertex and vertices of the paths it belongs to. We denote by u^i, $i \in [4]$, the vertices of $\mathbf{Im}(\phi)$ belonging to the boundary of $Q(u)$ and, for uniformity, also call them *joining vertices of $Q(u)$*.

Towards finding a model of H in the wall observe that the grid Γ contains as a subgraph a wall of height $(4\lambda + 1)\,(\mathbf{m}(H) + \mathbf{n}(H))$ such that each one of the boxes, either $\Gamma(v)$, $v \in V(H)$, or $Q(v)$, where v is the intersection of two paths in the image of ϕ contains a wall $W(v)$ of height $4 \deg(v)^2 + 1$ and the joining vertices of $\Gamma(v)$ (the vertices v^i, $i \in [4]$, respectively) belong to the perimeter of the wall and have distance at least 2 in it. Consider now the mapping of H to W where the boxes $\Gamma(v)$ and $Q(v)$ are mapped into subwalls $W(v)$ of W of height $4 \deg(v)^2 + 1$ joined together by vertex-disjoint paths as given by the orthogonal drawing ϕ. From Lemma 6, as every $W(v)$ has height $4 \deg(v)^2 + 1$ and its compass is embedded in a closed disk, there exist a vertex $z_v \in V(W(v))$ and $\deg(v)$ edge-disjoint paths from z_v to the joining vertices of $W(v)$. It is now easy to see that W contains a model of H. $\qquad \square$

Notice now that in the case when $\Delta(H) = O(1)$ we get the following.

Theorem 2. *There exists a computable function $f : \mathbb{N} \to \mathbb{N}$ such that for every multigraph G of Eüler genus γ and every connected graph H one of the following holds:*

1. *$\mathbf{tw}(G) \leq f(\gamma) \cdot \mathbf{n}(H)$,*

2. *G is not $\Delta(H)$-edge-connected,*

3. *$H \leq_{im} G$.*

The following two corollaries are immediate consequences of Theorems 1 and 2.

Corollary 1. *There exists a computable function $f : \mathbb{N} \to \mathbb{N}$ such that for every multigraph G of Eüler genus γ and every $k \in \mathbb{N}$ one of the following holds:*

1. $\mathbf{tw}(G) \leq f(\gamma) \cdot k^3$,

2. G *is not k-edge-connected,*

3. $K_{k+1} \leq_{im} G$.

Corollary 2. *There exists a computable function $f : \mathbb{N} \to \mathbb{N}$ such that for every multigraph G of Eüler genus γ and every $k \in \mathbb{N}$ one of the following holds:*

1. $\mathbf{tw}(G) \leq f(\gamma) \cdot k^2$,

2. G *is not 4-edge-connected,*

3. $(k \times k)$-*grid is an immersion of G.*

However, when H is the grid a straightforward argument gives the following result.

Theorem 3. *There exists a computable function $f : \mathbb{N} \to \mathbb{N}$ such that for every multigraph G that is embedded in a surface of Eüler genus γ and every $k \in \mathbb{N}$ one of the following holds:*

1. $\mathbf{tw}(G) \leq f(\gamma) \cdot k$.

2. G *is not 4-edge-connected.*

3. $(k \times k)$-*grid is an immersion of G.*

Proof. Let

$$f(\gamma, k) = 48 \cdot (\gamma + 1)^{\frac{3}{2}} \cdot ((4^3 + 3) \cdot k + 5).$$

Assume that G is 4-edge-connected and that $\mathbf{tw}(G) \geq f(\gamma, k)$. As $\mathbf{tw}(G) \geq f(\gamma, k)$, from Lemma 1 it follows that G contains as a subgraph a subdivided wall W of height $(4^3 + 3)k$, whose compass in G is embedded in a closed disk Δ.

Consider the k^2 subwalls of W of height $(4^3 + 1)$ that occur after removing from it the paths $P^{[v]}_{(4^3+3)j}$, $P^{[h]}_{(4^3+3)j}$, $i, j \in [k]$. For every $i, j \in [k]$, we denote by $W_{(i,j)}$ the subwall that is contained inside the disk that is defined by the paths $P^{(h)}_{(4^3+3)(i-1)}$, $P^{(h)}_{(4^3+3)i}$, $P^{[v]}_{(4^3+3)(j-1)}$, and $P^{[v]}_{(4^3+3)j}$. In the case where $j = 1$ and $i = 1$, we abuse notation and consider as $P^{(h)}_{(4^3+3)(j-1)}$ and $P^{[v]}_{(4^3+3)(j-1)}$ the paths $P^{(h)}_1$ and $P^{[v]}_1$, respectively.

From Lemma 6 and the hypothesis that G is 4-edge-connected, for $k = 4$, it follows that in the compass of each one of the subwalls $\{W_{(i,j)} \mid i, j \in [k]\}$ we may find a vertex $v_{(i,j)}$ and four edge-disjoint paths from $v_{(i,j)}$ to the vertices

$v_{(i,j)}^n$, $v_{(i,j)}^s$, $v_{(i,j)}^w$, and $v_{(i,j)}^e$, that lie in the northern, southern, western, and eastern path of the wall, respectively.

Finally, we consider the function $g((i,j)) = v_{(i,j)}$ that maps the vertex (i,j) of the $(k \times k)$-grid to the vertex $v_{(i,j)}$ of the wall $W_{(i,j)}$. Is now easy to see that g is an immersion model of the $(k \times k)$-grid in the compass of the wall W and the theorem follows as f is linear on k. $\qquad\square$

5 Conclusions

In this paper, we proved sufficient conditions for the containment of any connected graph H as an immersion in graphs of bounded genus. We would like to remark here that our proofs also hold if we, instead, consider the strong immersion relation where we additionally ask that the paths of the model f of H in G that correspond to the edges of H are internally disjoint from $f(V(H))$.

In our results, it appears that both big treewidth and the edge connectivity requirement are necessary in order to enforce the appearance of a graph as an immersion. A natural open problem to investigate is the existence of counterparts of our results for the case of the topological minor relation. Certainly, here edge connectivity should be replaced by vertex connectivity. However, what we can only report is that stronger conditions than just asking for sufficiently big treewidth are required for such an extension.

References

[1] F. N. Abu-Khzam and M. A. Langston. Graph coloring and the immersion order. In T. Warnow and B. Zhu, editors, *COCOON*, volume 2697 of *Lecture Notes in Computer Science*, pages 394–403. Springer, 2003.

[2] R. Belmonte, P. van 't Hof, M. Kaminski, D. Paulusma, and D. M. Thilikos. Characterizing graphs of small carving-width. In G. Lin, editor, *COCOA*, volume 7402 of *Lecture Notes in Computer Science*, pages 360–370. Springer, 2012.

[3] T. C. Biedl and M. Kaufmann. Area-efficient static and incremental graph drawings. In R. E. Burkard and G. J. Woeginger, editors, *ESA*, volume 1284 of *Lecture Notes in Computer Science*, pages 37–52. Springer, 1997.

[4] E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Subexponential parameterized algorithms on bounded-genus graphs and H-minor-free graphs. *J. ACM*, 52(6):866–893, 2005.

[5] M. DeVos, Z. Dvořák, J. Fox, J. McDonald, B. Mohar, and D. Scheide. Minimum degree condition forcing complete graph immersion. *ArXiv e-prints*, Jan. 2011.

[6] M. DeVos, K.-I. Kawarabayashi, B. Mohar, and H. Okamura. Immersing small complete graphs. *Ars Math. Contemp.*, 3(2):139–146, 2010.

[7] M. Ferrara, R. J. Gould, G. Tansey, and T. Whalen. On h-immersions. *Journal of Graph Theory*, 57(3):245–254, 2008.

[8] F. V. Fomin, P. A. Golovach, and D. M. Thilikos. Contraction obstructions for treewidth. *J. Comb. Theory, Ser. B*, 101(5):302–314, 2011.

[9] A. C. Giannopoulou, M. Kaminski, and D. M. Thilikos. Forbidding kuratowski graphs as immersions. *CoRR*, abs/1207.5329, 2012.

[10] A. C. Giannopoulou, I. Salem, and D. Zoros. Effective computation of immersion obstructions for unions of graph classes. In F. V. Fomin and P. Kaski, editors, *SWAT*, volume 7357 of *Lecture Notes in Computer Science*, pages 165–176. Springer, 2012.

[11] A. C. Giannopoulou and D. M. Thilikos. Optimizing the graph minors weak structure theorem. *CoRR*, abs/1102.5762, 2011.

[12] M. Grohe, K. ichi Kawarabayashi, D. Marx, and P. Wollan. Finding topological subgraphs is fixed-parameter tractable. In *STOC*, pages 479–488, 2011.

[13] K. ichi Kawarabayashi and Y. Kobayashi. List-coloring graphs without subdivisions and without immersions. In Y. Rabani, editor, *SODA*, pages 1425–1435. SIAM, 2012.

[14] A. V. Kostochka. Lower bound of the hadwiger number of graphs by their average degree. *Combinatorica*, 4(4):307–316, 1984.

[15] F. Lescure and H. Meyniel. On a problem upon configurations contained in graphs with given chromatic number. In L. D. Andersen, I. T. Jakobsen, C. Thomassen, B. Toft, and P. D. Vestergaard, editors, *Graph Theory in Memory of G.A. Dirac*, volume 41 of *Annals of Discrete Mathematics*, pages 325 – 331. Elsevier, 1988.

[16] B. Mohar and C. Thomassen. *Graphs on surfaces.* Johns Hopkins University Press, Baltimore, MD, 2001.

[17] N. Robertson, P. Seymour, and R. Thomas. Quickly excluding a planar graph. *J. Combin. Theory Ser. B*, 62(2):323–348, 1994.

[18] N. Robertson and P. D. Seymour. Graph minors. XIII. The disjoint paths problem. *J. Combin. Theory Ser. B*, 63(1):65–110, 1995.

[19] N. Robertson and P. D. Seymour. Graph minors. XVI. Excluding a nonplanar graph. *J. Comb. Theory, Ser. B*, 89(1):43–76, 2003.

[20] N. Robertson and P. D. Seymour. Graph minors XXIII. Nash-Williams' immersion conjecture. *J. Comb. Theory, Ser. B*, 100(2):181–205, 2010.

[21] A. Thomason. The extremal function for complete minors. *J. Combin. Theory Ser. B*, 81(2):318–338, 2001.

[22] P. Wollan. The structure of graphs not admitting a fixed immersion. *CoRR*, abs/1302.3867, 2013.