

Three-coloring triangle-free graphs on surfaces

VII. A linear-time algorithm*

Zdeněk Dvořák[†] Daniel Král'[‡] Robin Thomas[§]

Abstract

We give a linear-time algorithm to decide 3-colorability of a triangle-free graph embedded in a fixed surface, and a quadratic-time algorithm to output a 3-coloring in the affirmative case. The algorithms also allow to prescribe the coloring for a bounded number of vertices.

1 Introduction

This paper is the last part of a series aimed at studying the 3-colorability of graphs on a fixed surface that are either triangle-free, or have their triangles restricted in some way (throughout the paper, all colorings are proper, i.e., adjacent vertices have different colors). The main result of this paper is a linear-time algorithm to decide 3-colorability of a triangle-free graph embedded in a fixed surface. Embeddability in a surface is not a sufficient restriction by itself, as 3-colorability of planar graphs is NP-complete [7]. Restricting the triangles is natural in the light of the well-known theorem of Grötzsch stating that every planar triangle-free graph is 3-colorable.

A graph G is *4-critical* if every proper subgraph of G is 3-colorable, but G itself is not. Clearly, a graph is 3-colorable if and only if it has no 4-critical subgraph. As was shown by Thomassen [9] and later (with better bounds) by us [3], for every surface Σ , there are only finitely many 4-critical graphs of girth at least 5 that can be embedded in Σ . Hence, to decide whether a graph of girth at least 5 embeddable in Σ is 3-colorable, it suffices to test the presence of these finitely many subgraphs, which can be done in linear time using the algorithm of Eppstein [6].

*3 Jan 2015.

[†]Computer Science Institute (CSI) of Charles University, Malostranské náměstí 25, 118 00 Prague, Czech Republic. E-mail: rakdver@iuuk.mff.cuni.cz. Supported by project GA14-19503S (Graph coloring and structure) of Czech Science Foundation.

[‡]Warwick Mathematics Institute, DIMAP and Department of Computer Science, University of Warwick, Coventry CV4 7AL, United Kingdom. E-mail: D.Kral@warwick.ac.uk. The previous affiliation: Computer Science Institute (CSI) of Charles University.

[§]School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332. E-mail: thomas@math.gatech.edu. Partially supported by NSF Grants No. DMS-0739366 and DMS-1202640.

The situation is more complicated with triangle-free graphs. The Mycielski graph of any odd cycle embeds in any surface other than the sphere. Furthermore, Youngs [10] gave more general infinite families of 4-critical triangle-free graphs embeddable in any non-orientable surface. However, in [4] we showed that most of the faces of such a graph are of length 4. In order to state the result precisely, let us first give a few definitions.

A *surface* is a compact connected 2-manifold with (possibly null) boundary. Each component of the boundary is homeomorphic to a circle, and we call it a *cuff*. For non-negative integers a , b and c , let $\Sigma(a, b, c)$ denote the surface obtained from the sphere by adding a handles, b crosscaps and removing the interiors of c pairwise disjoint closed discs. The classification theorem of surfaces shows that every surface is homeomorphic to $\Sigma(a, b, c)$ for some choice of a , b and c . The *Euler genus* $g(\Sigma)$ of a surface Σ homeomorphic to $\Sigma(a, b, c)$ is defined as $2a + b$. Consider a graph G embedded in the surface Σ ; when useful, we identify G with the topological space consisting of the points corresponding to the vertices of G and the simple curves corresponding to the edges of G . A *face* f of G is a maximal connected subset of $\Sigma - G$. By the *length* $|f|$ of f , we mean the sum of the lengths of the boundary walks of f (in particular, if an edge appears twice in the boundary walks, it contributes 2 to $|f|$). A face f is *2-cell* if it is homeomorphic to an open disk, and it is *closed 2-cell* if additionally its boundary forms a cycle in G .

Finally, we are ready to state the result.

Theorem 1.1 ([4, Theorem 1.3]). *There exists a constant κ with the following property. Let G be a graph embedded in a surface of Euler genus g . Let t be the number of triangles in G and let c be the number of 4-cycles in G that do not bound a 2-cell face. If G is 4-critical, then*

$$\sum_{f \text{ face of } G} (|f| - 4) \leq \kappa(g + t + c - 1).$$

Furthermore, in the previous paper of the series [5], we designed a 3-coloring algorithm for graphs with almost all faces of length 4. A graph H is a *quadrangulation* of a surface Σ if all faces of H are closed 2-cell and have length 4 (in particular, the boundary of Σ is formed by a set of pairwise vertex-disjoint cycles in H , called the *boundary cycles* of H). A vertex of G contained in the boundary of Σ is called a *boundary vertex*.

Theorem 1.2 ([5, Theorem 1.3]). *For every surface Σ and integer k , there exists a linear-time algorithm with input*

- G : a quadrangulation of Σ with at most k boundary vertices, and
- ψ : a 3-coloring of the boundary cycles of G ,

which correctly decides whether there exists a 3-coloring φ of G such that $\varphi(v) = \psi(v)$ for every boundary vertex v of G . In the affirmative case, the algorithm also outputs such a coloring φ .

By combining Theorems 1.1 and 1.2, we obtain a straightforward algorithm to test 3-colorability of a triangle-free graph embedded in a fixed surface, at least under the assumption that all 4-cycles in G bound 2-cell faces—enumerate all subgraphs H of G such that

$$\sum_{h \text{ face of } H} (|h| - 4) \leq \kappa g \tag{1}$$

and test whether they are 3-colorable. It can be shown that there are at most $|V(G)|^{5\kappa g}$ subgraphs of G satisfying (1), and thus for any fixed surface, we obtain a polynomial-time algorithm. However, the exponent of the polynomial bounding the complexity of this algorithm depends on the surface. In this paper, we use a more involved argument to design a linear-time algorithm deciding 3-colorability. Furthermore, similarly to Theorem 1.2, we can allow a bounded number of precolored vertices in the considered graph. The following is our first main result.

Theorem 1.3. *For every surface Σ and integer k , there exists a linear-time algorithm with input*

- G : a triangle-free graph embedded in Σ with at most k boundary vertices, and
- ψ : a function from boundary vertices of G to $\{1, 2, 3\}$,

which correctly decides whether there exists a 3-coloring φ of G such that $\varphi(v) = \psi(v)$ for every boundary vertex v of G .

Note that unlike Theorem 1.2, Theorem 1.3 does not return a 3-coloring if one exists; indeed, the algorithm only decides whether there exists a critical subgraph. Let us remark that this is often the case; even in the case of planar graphs, a linear-time algorithm to actually find a 3-coloring guaranteed by Grötzsch’s theorem was only designed recently [1]. In order to report a 3-coloring, we need to go deeper into the proof of Theorem 1.1. The proof is based on the method of discharging and reducible configurations, and in Section 5, we convert it to the following algorithm.

Theorem 1.4. *For every surface Σ and integer k , there exists a quadratic-time algorithm with input*

- G : a triangle-free graph embedded in Σ with at most k boundary vertices, and
- ψ : a function from boundary vertices of G to $\{1, 2, 3\}$,

which correctly decides whether there exists a 3-coloring φ of G such that $\varphi(v) = \psi(v)$ for every boundary vertex v of G , and outputs such a coloring in the affirmative case.

Let us remark that we believe that there exists a linear-time algorithm to output a 3-coloring using ideas similar to those of Theorem 1.3; however, there are significant technical challenges in designing it and we leave this as an open problem. Let us also remark that in Theorems 1.3 and 1.4, it would suffice to only forbid the existence of contractible triangles, as we can deal with non-contractible ones by cutting the surface along them (see Theorem 4.3 for a similar idea used to eliminate non-contractible 4-cycles).

In the following section, we recall the results and definitions from the previous papers of the series we are going to need. In Section 3, we define a key notion of a free set of faces and apply it to the special case of graphs embedded in the disk. In Section 4, we give the linear-time decision algorithm. Finally, in Section 5, we give the algorithm to output a 3-coloring if one exists.

2 Definitions and previous results

We need a stronger form of Theorem 1.1 which deals with graphs with precolored cycles. First, let us give several definitions. Suppose that a graph G is embedded in a surface Σ so that every cuff of Σ traces a cycle in G , let H be a subgraph of G , and let h be a face of H . Let Σ_h denote the surface whose interior is homeomorphic to h , let $\theta_h : \Sigma_h \rightarrow \Sigma$ be a continuous function whose restriction to the interior of Σ_h is a homeomorphism to h , and let $G_h = \theta_h^{-1}(G)$. Let $s : \mathbf{Z}^+ \rightarrow \mathbf{R}$ be defined by

$$s(n) = \begin{cases} 0 & \text{if } n \leq 4 \\ 4/4113 & \text{if } n = 5 \\ 72/4113 & \text{if } n = 6 \\ 540/4113 & \text{if } n = 7 \\ 2184/4113 & \text{if } n = 8 \\ n - 8 & \text{if } n \geq 9. \end{cases}$$

To each 2-cell face f of G , we assign a weight $w_0(f) = s(|f|)$. If f is not 2-cell, then let $w_0(f) = |f|$. For a surface Π of Euler genus g with c cuffs, let $s(\Pi) = 6c - 6$ if $g = 0$ and $c \leq 2$, and $s(\Pi) = 120g + 48c - 120$ otherwise. For a real number η and a face f of G , let $w_\eta(f) = w_0(f) + \eta s(\Sigma_f)$. Let

$$w_\eta(G) = \sum_{f \text{ face of } G} w_\eta(f).$$

Theorem 2.1 ([5, Lemma 5.2]). *There exists a constant $\eta > 0$ such that the following holds. Let G be a triangle-free graph embedded in a surface Σ without non-contractible 4-cycles, so that every cuff of Σ traces a cycle in G , and let B be the union of boundary cycles of G . There exists a subgraph H of G such that $B \subseteq H$, $w_\eta(H) \leq w_\eta(B)$ and for every face h of H , every 3-coloring of the boundary of h extends to a 3-coloring of G_h .*

Throughout the rest of the paper, let η denote the constant of Theorem 2.1. We also need a stronger variant for the disk, see Corollary 5.4 in [4].

Theorem 2.2. *Let G be a triangle-free graph embedded in the disk with boundary cycle B . Then either every 3-coloring of B extends to a 3-coloring of G , or there exists a connected subgraph $H \subseteq G$ such that $B \subsetneq H$ and $w_0(H) \leq s(|B| - 2)$.*

Proof. If some 3-coloring φ of B does not extend to a 3-coloring of G , then let $H \supsetneq B$ be a minimal subgraph of G such that φ does not extend to a 3-coloring of H . By Grötzsch's theorem and the minimality of H , we conclude that H is connected. Furthermore, it is easy to see that H is B -critical (in the sense defined in Section 2 of [4]), and by [4, Corollary 5.4], we have $w_0(H) \leq s(|B| - 2)$. \square

Consider a graph G embedded in a surface Σ . A cycle K is *contractible* if there exists a closed disk $\Delta \subseteq \Sigma$ with boundary equal to K . For a cuff C , let $\Sigma + \widehat{C}$ denote the surface obtained from Σ by adding an open disk disjoint from Σ and with boundary equal to C ; we say that $\Sigma + \widehat{C}$ is obtained from Σ by *patching a cuff*. A cycle K *surrounds a cuff C* if K is not contractible in Σ , but it is contractible in $\Sigma + \widehat{C}$.

For algorithms, we need to specify how a graph G embedded in Σ is represented. Let g be the Euler genus of Σ and c the number of cuffs of Σ . We use the following variant of the polygonal representation. Let H be a graph drawn in Σ such that

- the boundary of every cuff traces a cycle in H ,
- every edge of H is either equal to an edge of G , or intersects G only in vertices,
- H has exactly one face and this face is homeomorphic to an open disk Λ , and
- H has at most $2(g + c - 1)$ vertices of degree three.

The graph G embedded in Σ is represented by a graph G' drawn in a closed disk Δ , such that there exists a continuous surjection $\theta : \Delta \rightarrow \Sigma$ satisfying

- $G' = \theta^{-1}(G)$,
- the restriction of θ to the interior of Δ is a homeomorphism to Λ , and
- θ maps the boundary of Δ to the boundary walk of the face Λ of H .

Note that there exist pairwise internally disjoint closed intervals $A_1, B_1, A_2, B_2, \dots, A_p, B_p$ for some $p \leq 3(g + c) + 1$ such that the restriction of θ to each of them is injective, $\theta(A_i) = \theta(B_i)$ for $1 \leq i \leq p$, and $\bigcup_{i=1}^p \theta(A_i)$ is the subgraph of H consisting of edges not contained in any of the boundary cycles of G . Hence, the surface Σ is obtained from the disk Δ by identifying A_i with

B_i for $1 \leq i \leq p$, in the direction prescribed by θ . We call this representation of an embedded graph a *normal representation*. Let us remark that a normal representation can be obtained from any other common representation of an embedded graph in linear time.

We also use a data structure we designed in [5], see Lemma 4.6.

Lemma 2.3. *For any integer $d \geq 0$ and every surface Σ , there exists a data structure as follows. The data structure represents a graph G with a 2-cell embedding in Σ and supports the following operations in amortized constant time (depending only on d and Σ):*

- (a) *Removal of an edge or an isolated vertex.*
- (b) *For any vertex $v \in V(G)$, deciding whether there exists a closed walk W of length at most d with $v \in V(W)$ such that W is not null-homotopic even after patching any one cuff of Σ with a disk, and finding such a walk if that is the case.*
- (c) *For any vertex $v \in V(G)$ and any set D of cuffs of Σ , letting Σ' be a surface obtained from Σ by patching all the cuffs in D and letting $\Lambda \subseteq \Sigma'$ be an open disk containing all the patches, deciding whether there exists a closed walk W in G of length at most d such that W contains v and is homotopically equivalent (in Σ) to the boundary of Λ , and finding such a walk if that is the case.*

Given a normal representation of G , the data structure can be initialized in $O(|V(G)|)$ time.

In [5], we designed several useful algorithms. One of them can be used to eliminate contractible (≤ 4)-cycles.

Lemma 2.4 ([5], Lemma 4.9). *For any surface Σ , there exists a linear-time algorithm that for any graph G with a 2-cell embedding in Σ returns its subgraph H such that*

- *H is 2-cell embedded in Σ and all boundary cycles of G belong to H ,*
- *all contractible cycles in H of length at most 4 bound 2-cell faces, and*
- *all vertices and edges of G that do not belong to H are drawn in 2-cell (≤ 4)-faces of H .*

We will need this result in combination with the algorithm of Dvořák, Kawarabayashi and Thomas [1].

Theorem 2.5. *There exists a linear-time algorithm as follows. Let G be a plane triangle-free graph with the outer face bounded by a cycle C of length at most 5. Given a 3-coloring ψ of C , the algorithm returns a 3-coloring of G that extends ψ .*

In particular, in the situation of Lemma 2.4, any 3-coloring of H can be extended to a 3-coloring of G in time $O(|V(G) \setminus V(H)|)$.

Consider a graph G embedded in a surface Σ . A subgraph H of G is *non-essential* if there exists $\Lambda \subset \Sigma$ containing H , where Λ is either an open disk, or an open disk with a hole whose boundary is equal to a cuff of Σ . A subgraph H of G is *essential* if it is not non-essential. We say that a surface Σ' is *at most as complex as* Σ if Σ' has smaller genus than Σ , or Σ' has the same genus and fewer cuffs than Σ , or Σ' is homeomorphic to Σ . For a graph G embedded in Σ , let $b(G)$ denote the multiset of the lengths of the boundary cycles of G . For two multisets S, T of integers such that $|S| = |T| = m$, we say that S *dominates* T if there exists an ordering s_1, \dots, s_m of the elements of S and an ordering t_1, \dots, t_m of the elements of T such that $s_i \geq t_i$ for $i = 1, \dots, m$.

The following algorithm is useful when dealing with essential subgraphs; let us recall that the notations Σ_h and G_h were defined at the beginning of this section.

Lemma 2.6 ([5], Lemma 4.7). *Let $\nu(\Sigma, k)$ be any function. For any surface Σ and integer $k \geq 0$, there exists a constant σ and a linear-time algorithm as follows. Let G be a graph 2-cell embedded in Σ with boundary cycles B_1, \dots, B_c of total length at most k . The algorithm returns a subgraph H of G with at most σ vertices such that $B_1 \cup \dots \cup B_c \subseteq H$ and for each face h of H , G_h (in its embedding in Σ_h) does not contain any connected essential subgraph with fewer than $\nu(\Sigma_h, k_h)$ edges, where k_h is the sum of the lengths of the boundary cycles of G_h . Furthermore, Σ_h is at most as complex as Σ , and if Σ_h is homeomorphic to Σ , then $b(G)$ dominates $b(G_h)$.*

We also need a similar algorithm to deal with the cylinder case.

Lemma 2.7 ([5], Lemma 4.3). *Let d be a positive integer. There exists a linear-time algorithm that, given a graph G that is 2-cell embedded in the cylinder Σ with boundary cycles B_1 and B_2 of length at most d , returns a sequence C_0, C_1, \dots, C_m of non-contractible cycles of G of length at most d such that*

- $C_0 = B_1$ and $C_m = B_2$,
- for $0 \leq i < m$, the cycle C_i is contained in the part of Σ between B_1 and C_{i+1} , and
- either C_i intersects C_{i+1} , or the subcylinder of Σ between C_i and C_{i+1} contains no non-contractible cycle of length at most d distinct from C_i and C_{i+1} .

3 Freedom

Let G be a triangle-free graph with a 2-cell embedding in a surface Σ . Let S be a set of faces of G and let W be a contractible closed walk in G forming the

boundary of an open disk $\Lambda \subset \Sigma$. We say that W binds S (with respect to Λ) if $\Lambda \notin S$ and

$$\sum_{f \in S, f \subseteq \Lambda} w_0(f) \geq s(|W|).$$

We say that S is k -free if no closed walk of length at most k binds S . The key observation is that the presence of a large free set ensures the possibility to extend a precoloring.

Lemma 3.1. *Let G be a triangle-free graph with a 2-cell embedding in the disk Δ , with boundary cycle B . If G contains a $(|B| - 2)$ -free set S of faces such that $\sum_{f \in S} w_0(f) > s(|B| - 2)$, then every 3-coloring of B extends to a 3-coloring of G .*

Proof. Suppose for a contradiction that some 3-coloring of B does not extend to a 3-coloring of G . By Theorem 2.2, there exists a connected $H \subseteq G$ with $B \subsetneq H$ such that $w_0(H) \leq s(|B| - 2)$. Note that $|h| \leq |B| - 2$ for every face h of H . Since S is $(|B| - 2)$ -free, we have

$$\sum_{f \in S, f \subseteq h} w_0(f) \leq w_0(h),$$

where we only write the non-strict inequality since h may belong to S . Therefore,

$$\sum_{f \in S} w_0(f) \leq \sum_{h \text{ face of } H} w_0(h) = w_0(H) \leq s(|B| - 2),$$

which is a contradiction. \square

We will need the following consequence. If G is a graph embedded in Σ and Λ is an open disk whose boundary is contained in G , then let $G - \Lambda$ be the graph obtained from G by removing all vertices and edges drawn in Λ .

Corollary 3.2. *Let G be a triangle-free graph with a 2-cell embedding in a surface Σ , let W be a closed walk in G bounding an open disk $\Lambda \subset \Sigma$, and let S be a set of faces such that $\sum_{f \in S, f \subseteq \Lambda} w_0(f) > s(|W| - 2)$. If S is $(|W| - 2)$ -free, then every 3-coloring of $G - \Lambda$ extends to a 3-coloring of G .*

Let us design an algorithm enabling us to perform this reduction efficiently. We need an observation enabling us to simplify closed walks of a given homotopy. By *drilling a hole* in a face f of an embedded graph, we mean deleting the interior of an arbitrary closed disk contained in f (so we obtain a new cuff disjoint from the boundary of f).

Lemma 3.3. *Let G be a graph with a 2-cell embedding in a surface Σ , let S be a set of faces of G and let Σ' be the surface obtained from Σ by drilling holes in the faces of S . Let $\Lambda \subseteq \Sigma$ be an open disk such that $\Sigma \setminus \Sigma' \subseteq \Lambda$. Let W be a closed walk in G homotopically equivalent in Σ' to the boundary of Λ , and let H_W be the subgraph of G consisting of the vertices and the edges of W . Let T be the set of faces of H_W in its embedding in Σ that intersect $\Sigma \setminus \Sigma'$.*

If G (in Σ) contains no connected essential subgraph with at most $|W|$ edges, then each face of T is 2-cell and $\sum_{f \in T} |f| \leq |W|$.

Proof. Since G contains no connected essential subgraph with at most $|W|$ edges, we conclude that there exists $\Lambda_H \subseteq \Sigma$ containing H_W such that Λ_H is either an open disk, or an open disk with a hole whose boundary is equal to a cuff C of Σ . In the latter case, let Λ'_H be the open disk obtained from Λ_H by patching the hole corresponding to C . In the former case, let $\Lambda'_H = \Lambda_H$.

Consider H_W as embedded in Λ'_H . Let f_0 be the face of H_W containing the boundary of Λ'_H , and if $\Lambda'_H \neq \Lambda_H$, then let f_1 be the face of H_W containing $\Lambda'_H \setminus \Lambda_H$. Since H_W is connected, all faces of H_W except for f_0 are 2-cell. Since W is homotopically equivalent to the boundary of Λ , we have $f_0, f_1 \notin T$. Consequently, all faces of T are 2-cell. Furthermore, if two faces of T share an edge e , then e appears in W at least twice (since W has the same winding number ± 1 around both of the faces). Therefore, $\sum_{f \in T} |f| \leq |W|$. \square

We can now design the following subroutine.

Lemma 3.4. *For any integer $k \geq 0$ and a surface Σ , there exists a linear-time algorithm as follows. Let G be a triangle-free graph with a 2-cell embedding in a surface Σ such that every connected essential subgraph of G has at least $k + 1$ edges, and let f be a face of G . The algorithm decides whether $\{f\}$ is k -free, and if not, returns a closed walk W of length at most k that binds $\{f\}$ with respect to an open disk $\Lambda \subseteq \Sigma$ such that among all such walks, $|W|$ is minimal, and additionally $\{\Lambda\}$ is k -free in $G - \Lambda$.*

Proof. Let Σ' be the surface obtained from Σ by drilling a hole inside f , and let Λ' be the open disk removed from Σ in order to create Σ' . Build the data structure of Lemma 2.3 for G with $d = k$, and find (in linear time) a shortest walk W_0 of length at most k such that W_0 is homotopically equivalent in Σ' to the boundary of Λ' . If no such walk exists, then $\{f\}$ is k -free.

Otherwise, let $t = |W_0|$. Let H_0 be the subgraph of G consisting of the vertices and edges of G contained in W_0 , and let Λ_0 be the face of H_0 containing f . By Lemma 3.3, Λ_0 is an open disk bounded by a closed walk W'_0 of length at most t . By the minimality of $|W_0|$, it follows that $|W'_0| = t$.

We now proceed as follows. Set $W = W'_0$ and $\Lambda = \Lambda_0$, and remove in the data structure all vertices and edges drawn in Λ so that the data structure now represents $G - \Lambda$. We process all remaining vertices not contained in W in order, and for each of them test in constant time whether there exists a closed walk of length t passing through it and homotopically equivalent to the boundary of Λ' . Whenever we find such a walk, we repeat the procedure of the previous paragraph and replace W and Λ by the obtained walk and the open disk bounded by it, and remove the vertices and edges so that the data structure represents $G - \Lambda$.

At the end of this procedure, we end up with a closed walk W and an open disk Λ satisfying the conclusions of the lemma (if $\Lambda = f$, then $\{f\}$ is k -free). \square

Let us remark that in the situation of Lemma 3.4, if $\{f\}$ is not k -free, then the reduction of Corollary 3.2 applies ($\{f\}$ is $(|W| - 2)$ -free since $|W|$ is minimal). Another subroutine is used to search for a binding walk around several faces.

Lemma 3.5. *For all integers $k, b \geq 0$ and a surface Σ , there exists a linear-time algorithm as follows. Let G be a triangle-free graph with a 2-cell embedding in a surface Σ such that every connected essential subgraph of G has at least $k + 1$ edges, and let S be a set of faces of G such that $|S| \leq b$. The algorithm decides whether S is k -free, and if not, returns a shortest closed walk W that binds S .*

Proof. Using the algorithm of Lemma 3.4, we test whether single-element subsets of S are k -free, and record the obtained shortest walks binding them (if any).

Now, for every $D \subseteq S$ such that $|D| \geq 2$, let $t \leq k$ be the largest integer such that $\sum_{f \in D} w_0(f) \geq s(t)$, and proceed as follows. Let Σ' be the surface obtained from Σ by drilling a hole inside each face of D . Let Λ be an open disk in Σ containing $\Sigma \setminus \Sigma'$. Build the data structure of Lemma 2.3 for G with $d = t$, and determine whether there exists a closed walk W_0 of length at most t that is homotopically equivalent in Σ' to the boundary of Λ .

Suppose that we found such a walk W_0 , chosen to be the shortest possible. Let H_{W_0} be the subgraph of G consisting of the vertices and edges of W_0 and let T be the set of faces of H_{W_0} intersecting $\Sigma \setminus \Sigma'$. By Lemma 3.3, we have

$$\sum_{f \in D} w_0(f) \geq s(t) \geq s\left(\sum_{h \in T} |h|\right) \geq \sum_{h \in T} s(h),$$

and either $|T| = 1$ or the last inequality is strict. Hence, there exists $h \in T$ bounded by a closed walk W such that $\sum_{f \in D, f \subseteq h} w_0(f) \geq w(h)$, and either the inequality is strict, or $\bigcup D \subseteq h$. Therefore, W binds S , and we record W .

In the end, we return a shortest recorded walk that binds S (or that S is k -free if no walk was recorded). \square

Again, if S is not k -free, then the minimality of $|W|$ ensures that in the situation of Lemma 3.5, S is $(|W| - 2)$ -free. We can now describe the reduction algorithm.

Lemma 3.6. *For an integer $k \geq 4$, a rational number $r \geq 0$ and a surface Σ , there exists a linear-time algorithm as follows. Let G be a triangle-free graph with a 2-cell embedding in a surface Σ such that every connected essential subgraph of G has at least $k + 1$ edges. The algorithm returns $G' \subseteq G$ whose embedding in Σ induced by the embedding of G is 2-cell such that*

- every boundary cycle of G belongs to G' ,
- every 3-coloring of G' extends to a 3-coloring of G , and
- either $w_0(G') \leq r$, or G' contains a k -free set of faces S such that

$$\sum_{f \in S} w_0(f) > r.$$

Proof. Let $b = \lceil r/s(5) \rceil + 1$. Note that there exist only finitely many multisets M of integers greater than 4 such that $\sum_{t \in M} s(t) \leq r$. Let m_0 denote the number of such multisets.

We construct a sequence $(G_0, S_0), (G_1, S_1), \dots, (G_m, S_m)$ such that for $i = 0, \dots, m$, G_i is a subgraph of G , S_i is a set of faces of G_i , and

- the embedding of G_i in Σ is 2-cell and every boundary cycle of G belongs to G_i ,
- every 3-coloring of G_i extends to a 3-coloring of G ,
- $|S_i| \leq b$, and
- for every $f \in S_i$, $|f| \geq 5$ and the set $\{f\}$ is k -free in G_i .

Let G_0 be the subgraph of G obtained by applying the algorithm of Lemma 2.4 and suppressing all faces of length two, and let $S_0 = \emptyset$. Assuming we already constructed (G_i, S_i) for some $i \geq 0$, we proceed as follows.

We apply the algorithm of Lemma 3.5 to test whether S_i is k -free in G_i . If not, let W_0 be the closed walk returned by the algorithm that binds S_i with respect to an open disk $\Lambda_0 \subset \Sigma$. By Corollary 3.2, every 3-coloring of $G_i - \Lambda_0$ extends to a 3-coloring of G_i , and thus also to a 3-coloring of G . Apply the algorithm of Lemma 3.4 to $G_i - \Lambda_0$ and its face Λ_0 , and let W be the closed walk returned by the algorithm that binds $\{\Lambda_0\}$ with respect to an open disk $\Lambda \subset \Sigma$, or set $W = W_0$ and $\Lambda = \Lambda_0$ when $\{\Lambda_0\}$ already is k -free. Let $G_{i+1} = G_i - \Lambda$ and let S_{i+1} be obtained from S_i by removing the faces contained in Λ and by adding Λ . Note that by Corollary 3.2, every 3-coloring of G_{i+1} extends to a 3-coloring of $G_i - \Lambda_0$, and thus also to a 3-coloring of G . Furthermore, $|S_{i+1}| \leq |S_i| \leq b$.

Hence, assume that S_i is k -free in G_i . If either $\sum_{f \in S_i} w_0(f) > r$ or S_i contains all faces of G_i of length at least 5, then we set $m = i$ and end the procedure. Otherwise, note that $|S_i| \leq r/s(5) \leq b - 1$, and let $f \notin S$ be a face of G_i of length at least 5. Let $\Lambda \subset \Sigma$ be an open disk found using Lemma 3.4 applied for f such that $f \subseteq \Lambda$ and Λ is k -free in $G_i - \Lambda$. By the choice of G_0 , the boundary walk of Λ has length at least 5. Let $G_{i+1} = G_i - \Lambda$ and let S_{i+1} be obtained from S_i by removing the faces contained in Λ and by adding Λ . Note that by Corollary 3.2, every 3-coloring of G_{i+1} extends to a 3-coloring of G_i , and thus also to a 3-coloring of G .

This finishes the description of the construction of the sequence G_0, G_1, \dots, G_m . Let us now give a bound on the length of the sequence. For $i \geq 0$, let M_i denote the sequence of the lengths of faces in S_i in the non-increasing order. Observe that in order to obtain S_{i+1} , some subset X of faces in S_i is replaced by another face Λ , and every face in X is strictly shorter than Λ by the assumption that $\{f\}$ is k -free in G_i for every $f \in S_i$. Hence, M_{i+1} is lexicographically strictly larger than M_i . Consequently, $M_i \neq M_j$ for every $i \neq j$, and thus there exist at most m_0 values of i such that $\sum_{f \in S_i} w_0(f) \leq r$. Furthermore, observe that if S_i is not k -free, then $|S_{i+1}| < |S_i|$, and thus there do not exist b consecutive values of i such that S_i is not k -free. We conclude that the described

algorithm terminates after $m \leq b(m_0 + 1)$ steps, which is a constant depending only on r .

By the construction, either S_m is k -free in G_m and $\sum_{f \in S_m} w_0(f) > r$, or $w_0(G_m) \leq r$. Therefore, we can set $G' = G_m$. \square

As a corollary, we obtain the special case of Theorem 1.3 when Σ is the disk.

Corollary 3.7. *For every integer n , there exists a linear-time algorithm with input*

- G : a triangle-free graph embedded in the disk Δ with the boundary cycle B of length at most n , and
- ψ : a 3-coloring of B ,

which either correctly decides that there exists a 3-coloring φ of G extending ψ , or returns a subgraph $G' \subseteq G$ such that $B \subseteq G'$, $w_\eta(G') \leq s(|B| - 2)$ and ψ does not extend to a 3-coloring of G' .

Proof. Let G' be the subgraph of G obtained by applying the algorithm of Lemma 3.6 with $k = |B| - 2$ and $r = s(|B| - 2)$. Note that ψ extends to a 3-coloring of G if and only if it extends to a 3-coloring of G' . If $w_0(G') \leq r = s(|B| - 2)$, then we can decide whether ψ extends to a 3-coloring of G' using the algorithm of Theorem 1.2, by testing all the possible colorings of the vertices incident with faces of G' of length greater than 4 (there are at most $5s(|B| - 2)/s(5)$ such vertices). Otherwise, G' contains a $(|B| - 2)$ -free set of total weight greater than $s(|B| - 2)$, and Lemma 3.1 implies that ψ extends to a 3-coloring of G' . \square

4 Linear-time decision algorithm

The algorithm for general surfaces is now only mildly more complicated. Let us prepare by giving its special case where the graph has no non-contractible 4-cycles or small connected essential subgraphs.

Theorem 4.1. *For every surface Σ and integer k , there exists a linear-time algorithm with input*

- G : a triangle-free graph 2-cell embedded in Σ with boundary cycles B , without non-contractible 4-cycles, and without connected essential subgraphs with at most $w_\eta(B)$ edges, such that $|V(B)| \leq k$, and
- ψ : a 3-coloring of B ,

which either correctly decides that there exists a 3-coloring φ of G extending ψ , or returns a subgraph $G' \subseteq G$ such that $B \subseteq G'$, $w_\eta(G') \leq w_\eta(B)$ and ψ does not extend to a 3-coloring of G .

Proof. We proceed by induction on the number of cuffs of Σ ; hence, we can assume that the algorithm exists for all surfaces obtained from Σ by patching at least one cuff. Firstly, for each cuff C of Σ , run the algorithm for the embedding of G in the surface obtained from Σ by patching C and for the restriction of ψ to $B - V(C)$. Suppose that the coloring does not extend, and thus we obtain a subgraph $G'' \subseteq G$ with $w_\eta(G'') \leq w_\eta(B - V(C))$ such that the restriction of ψ to $B - V(C)$ does not extend to a 3-coloring of G'' . Then, it suffices to set $G' = G'' \cup C$ and observe that $w_\eta(G') \leq w_\eta(B)$ and that ψ does not extend to a 3-coloring of G .

Hence, we can assume that for every cuff C , the restriction of ψ to $B - V(C)$ extends to a 3-coloring of G . Consequently, if $H \supseteq B$ is a subgraph of G such that ψ does not extend to a 3-coloring of H , then H is connected.

Now, apply the algorithm of Lemma 3.6 with $k = r = w_\eta(B)$, and let G' be the resulting subgraph. Let us first consider the case that $w_0(G') \leq r$. We can decide whether ψ extends to a 3-coloring of G' using the algorithm of Theorem 1.2, by testing all the possible colorings of the vertices incident with faces of G' of length greater than 4 (there are at most $5w_\eta(B)/s(5)$ such vertices). If it does, then ψ also extends to a 3-coloring of G . If ψ does not extend to a 3-coloring of G' , then G' is the subgraph as required in the conclusion of the theorem.

Next, consider the case that G' contains a k -free set S of faces of total weight greater than $w_\eta(B)$. Let H be the subgraph of G as in Theorem 2.1. Since $w_\eta(H) \leq w_\eta(B) = k$, every face h of H has length at most k . If the embedding of H is 2-cell, then $\sum_{f \in S, f \subseteq h} w_0(f) \leq w_0(h)$ since S is k -free. However, that would imply $\sum_{f \in S} w_0(f) \leq w_0(H) \leq w_\eta(B)$, which contradicts the choice of S .

Therefore, H contains a face that is not 2-cell. Since all faces of H have length at most k and H does not contain any connected essential subgraph with at most k edges, it follows that H is not connected. However, as we ensured before, this implies that ψ extends to a 3-coloring of H , and consequently to a 3-coloring of G' and of G . \square

Easily, we can allow the boundary cycles to have length 4 by a minor modification to the algorithm of Theorem 4.1: before running the algorithm, subdivide an edge in each cuff of length 4 by a vertex and extend ψ by giving the new vertex a color distinct from the colors of its neighbors.

Corollary 4.2. *For every surface Σ and integer k , there exists an integer $N_{\Sigma,k}$ and a linear-time algorithm with input*

- G : a triangle-free graph 2-cell embedded in Σ , without non-contractible 4-cycles other than the boundary cycles, and without connected essential subgraphs with at most $N_{\Sigma,k}$ edges, such that the union B of the boundary cycles of G has at most k vertices, and
- ψ : a 3-coloring of B ,

which correctly decides whether there exists a 3-coloring φ of G extending ψ .

Next, let us deal with other non-contractible 4-cycles.

Theorem 4.3. *For every surface Σ and integer k , letting $N_{\Sigma,k}$ be as in Corollary 4.2, there exists a linear-time algorithm with input*

- G : a triangle-free graph 2-cell embedded in Σ without connected essential subgraphs with at most $N_{\Sigma,k} + 8$ edges, such that the union B of the boundary cycles of G has at most k vertices, and
- ψ : a 3-coloring of B ,

which correctly decides whether there exists a 3-coloring φ of G extending ψ . Furthermore, if Σ is the cylinder (the sphere with two holes), then we can omit the restriction on the connected essential subgraphs.

Proof. By Corollary 3.7, we can assume that either Σ has non-zero genus or at least two cuffs.

Suppose first that Σ is the cylinder. Apply the algorithm of Lemma 2.7 with $d = k$, and let C_0, C_1, \dots, C_m be the resulting cycles. For $i = 1, \dots, m$, we can decide which 3-colorings of $C_{i-1} \cup C_i$ extend to the subgraph G_i of G drawn between $C_{i-1} \cup C_i$, as follows. If the distance between C_{i-1} and C_i is at most $N_{\Sigma,2k}$, then let P be a shortest path between C_{i-1} and C_i . Using Corollary 3.7 for the subgraph(s) of G contained in the 2-cell face(s) of $C_{i-1} \cup P \cup C_i$, we can decide which 3-colorings of $C_{i-1} \cup P \cup C_i$ extend to a 3-coloring of G_i , and thus also which 3-colorings of $C_{i-1} \cup C_i$ extend to a 3-coloring of G_i . If the distance between C_{i-1} and C_i is greater than $N_{\Sigma,2k}$, then since Σ is a cylinder, we conclude that G_i does not contain any connected essential subgraph with at most $N_{\Sigma,2k}$ vertices, and thus which 3-colorings of $C_{i-1} \cup C_i$ extend to G_i can be decided using the algorithm of Corollary 4.2.

Finally, we can combine the information by straightforward dynamic programming to determine whether ψ extends to a 3-coloring of G .

Hence, we can assume that Σ is not the cylinder. For each cuff C of Σ , let Σ_C be the surface obtained from Σ by patching C , and let f_C be the face of G bounded by C in its embedding in Σ_C . Let W_C be the closed walk obtained by applying the algorithm of Lemma 3.4 to $\{f_C\}$ with $k = 4$, and let Λ_C be the open disk bounded by W_C . Note that for any distinct cuffs C_1 and C_2 , the closures of Λ_{C_1} and Λ_{C_2} are disjoint, since every connected essential subgraph of G has more than 8 edges, and since Σ is not the cylinder.

Let C_1, \dots, C_t be the cuffs of Σ , let $\Sigma' = \Sigma \setminus (\Lambda_{C_1} \cup \dots \cup \Lambda_{C_t})$ and let G' be the subgraph of G embedded in Σ' . Note that Σ' is homeomorphic to Σ , that G' does not contain any non-contractible 4-cycles other than the boundary cycles, and that G' does not contain any connected essential subgraphs with at most $N_{\Sigma,k}$ edges. By Corollary 4.2, we can decide which 3-colorings of the boundary cycles $W_1 \cup \dots \cup W_t$ of G' extend to a 3-coloring of G' . For $i = 1, \dots, t$, we can decide which 3-colorings of $C_i \cup W_i$ extend to the subgraph of G drawn between C_i and W_i , by the cylinder case when W_i and C_i are vertex-disjoint, and by Corollary 3.7 otherwise. By combining this information, we can decide whether ψ extends to a 3-coloring of G . \square

To give the full algorithm, it now suffices to deal with the essential subgraphs.

Proof of Theorem 1.3. Without loss of generality, we can assume that the embedding of G in Σ is 2-cell.

Let $\nu(\Pi, t) = N_{\Pi, t}$ for every surface Π and integer t , where $N_{\Pi, t}$ is as in Corollary 4.2. Apply the algorithm of Lemma 2.6 to G , obtaining a subgraph H of G . For every face h of H , determine which 3-colorings of the boundary of h extend to a 3-coloring of G_h by applying the algorithm of Theorem 4.3. By combining this information, we can determine which 3-colorings of H extend to a 3-coloring of G , and thus also whether ψ extends to a 3-coloring of G . \square

5 Finding a 3-coloring

Note that the argument of Lemma 3.1 gives no way to find the 3-coloring that it proves to exist. Hence, to get such a 3-coloring, we need to delve deeper into the proof of Theorem 1.1.

The basic case we need to consider is that of graphs of girth at least 5. Note that in this case for every face h , $w_\eta(h) \geq \frac{s(5)}{5}|h|$, and thus $|E(H)| \leq \frac{5}{s(5)}w_\eta(H)$ for every graph H of girth at least 5 embedded in a surface. In particular, the size of the subgraph H in Theorem 2.1 becomes bounded by a constant dependent only on Σ and the lengths of the boundary cycles.

Let Σ be a surface and let G be embedded in Σ so that every (≤ 4) -cycle is one of the boundary cycles of G . Let B be the union of the boundary cycles of G . We say that G is *near-planar* if there exists a cuff C of Σ and an open subset Λ of Σ equal to an open disk with a hole bounded by C such that G consists of a component contained in Λ (which we call the *main component* of G) and of boundary cycles. For a near-planar graph G , let $m(G)$ denote the maximum size of the face of the main component of G in its embedding in Λ . We say that a near-planar graph G is *very exceptional* if either there exists an edge e with both ends in C such that $C + e$ in its embedding in Λ has a 2-cell face of length at most 6, or there exists a vertex v with neighbors $v_1, v_2, v_3 \in V(C)$ such that $C + \{vv_1, vv_2, vv_3\}$ in its embedding in Λ has two 2-cell faces of length 5. We need the following bound on the size of critical graphs of girth 5 in the disk, see Theorem 8.5 of [2].

Theorem 5.1. *There exists a constant α as follows. Let G be a graph of girth at least 5 embedded in the disk with boundary cycle B . Suppose that G does not contain a very exceptional subgraph. If some 3-coloring ψ of B does not extend to a 3-coloring of G , then there exists a subgraph $H \subseteq G$ such that $B \subset H$, ψ does not extend to a 3-coloring of H , $|V(H)| \leq \alpha|B|$ and every face of H has length at most $|B| - 5$.*

In particular, if $|B| \leq 9$ this implies that every 3-coloring of B extends to a 3-coloring of G . This was proved before by Thomassen [8], and closer analysis of his argument gives a quadratic algorithm to find such a 3-coloring of G .

Theorem 5.2. *There exists a quadratic-time algorithm with input*

- G : a graph of girth at least 5 embedded in the disk with boundary cycle B of length at most 9, such that B is an induced cycle and no vertex of G has more than two neighbors in B , and
- ψ : a 3-coloring of B ,

which outputs a 3-coloring φ of G such that $\varphi(v) = \psi(v)$ for every $v \in V(B)$.

Let G be a graph embedded in a surface without contractible cycles of length less than 5, and let C be a contractible cycle in G bounding an open disk Λ . We say that C is *clearable* if $|C| \leq 9$ and the subgraph of G drawn in the closed disk bounded by C does not contain a very exceptional subgraph.

The arguments proving all the mentioned results are based on the method of reducible configurations. Let G be a triangle-free graph 2-cell embedded in a surface Σ . Suppose that every 4-cycle of G is contained in the union B of the boundary cycles of G , and let ψ be a 3-coloring of B . By a *reducible configuration* in G we mean a connected subgraph F of G of bounded size, such that there exists a graph G' (which we call the *reduction with respect to ψ*) satisfying the following conditions.

- G' is a graph obtained from G by possibly adding non-crossing edges between vertices of F into some faces of G , by contracting some edges between vertices of F in the resulting graph, and by removing some vertices of F (thus, G' is a minor of a graph embedded in Σ , and thus it is also embedded in Σ).
- B is a subgraph of G' (i.e., no two vertices of B were identified in the reduction, and no vertex or edge of B was removed).
- Every 3-coloring of G' whose restriction to B is equal to ψ can be extended to a 3-coloring of G with the same restriction to B , in constant time.
- G' is triangle-free, and every 4-cycle of G' is contained in B .

We say that a subgraph of G' *touches* F if it contains either a vertex created by contractions during the reduction of F , or an edge added during the reduction of F .

Example 5.3. Suppose $F \subseteq G$ is one of the following:

- a vertex of degree at most 2 not contained in B , or
- an even cycle of vertices of degree three that is vertex-disjoint from B , or
- an odd cycle C of vertices of degree three that is vertex-disjoint from B and two adjacent vertices $x, y \in V(G) \setminus V(C)$, such that each of them has a neighbor in C .

In the first two cases, the reduction G' is equal to $G - V(F)$. In the last case, $G' = G - V(C)$. It is easy to check that every 3-coloring of G' extends to a 3-coloring of G .

The key result of [2] was showing the existence of reducible configurations in large enough graphs of girth 5. The following is a reformulation of Lemmas 5.10, 6.1, 6.2, 7.1 and 7.3 of [2].

Theorem 5.4. *For every surface Σ and an integer k , there exists a constant $R_{\Sigma,k}$ and a linear-time algorithm whose input is a triangle-free graph G with closed 2-cell embedding in a surface Σ such that*

- (a) *the union B of the boundary cycles of G has at most k vertices and $|V(G)| \geq R_{\Sigma,k}$,*
- (b) *every 4-cycle and every non-contractible cycle of length at most 7 is contained in B ,*
- (c) *every connected essential subgraph of G has at least 5 edges,*
- (d) *if a path P of length at most 4 has both ends in a boundary cycle K , then the distance between the ends of P in K is at most $|P|$, and*
- (e) *no two vertices of degree two of G contained in B are adjacent,*

and a 3-coloring ψ of B . The algorithm outputs either a clearable cycle in G , or a reducible configuration F with at most 17 vertices and the corresponding reduction G' with respect to ψ .

In the latter case, the reduction G' has the following property. Suppose that we are given a near-planar subgraph H' of G' such that $B \subseteq H'$, the main component of H' is 2-connected, only the main component of H' touches F , and every vertex of H' not belonging to B has degree at least three. Then,

- *we can in linear time construct a near-planar subgraph H of G such that $B \subseteq H$, $|V(H)| \leq |V(H')| + 10$ and $m(H) \leq m(H') + 5$, and*
- *H' is not very exceptional.*

Proof. Note that in this proof, we extensively refer to the concepts defined in [2]; we do not restate their definitions here, as the proof can only be understood in the context of [2] anyway.

In order to use the results [2], we need to consider conditions (I0), ..., (I9) for G . Note that (d) together with (b) and (c) implies that every vertex of $V(G) \setminus V(B)$ has at most one neighbor in B , that the boundary cycles are induced, and that the distance between any two boundary cycles is at least 5, which implies (I4) and (I7).

If G did not satisfy (I0), (I1), or (I2), it would contain one of the reducible configurations described in Example 5.3, and these reducible configurations trivially satisfy the conclusion of Theorem 5.4 (with $H = H'$, which is not very exceptional by (I4)). Note that conditions (I1) and (I2) do not specify an upper bound on the lengths of the cycles, but the inspection of their usage shows that they are only needed for cycles of length at most 9. In particular, the presence of these configurations can be tested in linear time, since any graph contains at most a linear number of (≤ 9) cycles consisting of vertices of degree three.

Condition (I3) is implied by the assumption that the embedding of G is closed 2-cell and by (b). Condition (I5) is implied by (I0) and (e). Condition (I6) is implied by the assumption that the embedding of G is closed 2-cell. Condition (I8) is implied by (b). Finally, if condition (I9) is not satisfied, then G contains a clearable cycle, which can be detected in linear time using the ideas of [5, Lemma 4.9].

Hence, we can assume that G satisfies conditions (I0)–(I9). Together with (c) and (d), this implies that G is well-behaved.

Let $M = B$, let g be the genus of Σ , let c be the number of cuffs of Σ , and let $R_{\Sigma,k} = 8g + 8c + 9|E(B)|$. By [2, Lemma 5.10] and (a), a configuration isomorphic to one of R1, \dots , R7 appears in G and does not touch (in the sense of [2]) $M = B$. By [2, Lemma 6.1] and (a), a good configuration γ strongly appears in G . We can find such a configuration (in fact, all of them) in linear time, because every good configuration has a connected subgraph with at most 12 vertices, each of them of degree at most four in G , such that every vertex of the configuration has a neighbor in the subgraph.

Let G' be the γ -reduction of G with respect to ψ . By [2, Lemma 6.2], G' is triangle-free and every 4-cycle of G' is contained in B , and thus γ forms a reducible configuration (with at most 17 vertices involved in the reduction, giving the bound on the size of F). Let H' be a near-planar subgraph H' of G' such that $B \subseteq H'$, the main component of H' is 2-connected, only the main component of H' touches F , and every vertex of H' not belonging to B has degree at least three. By [2, Lemma 7.3], H' is not very exceptional.

At the beginning of Section 7 of [2], we assign to each face f'' of H' (playing the role of G'' in that section) a subgraph $J_{f''}$ of G . Let H be the subgraph of G equal to the union of such subgraphs $J_{f''}$ over all faces f'' of H' ; clearly, this subgraph can be constructed in linear time. By [2, Lemma 7.1], we have $|V(H)| \leq |V(H')| + 10$ and $m(H) \leq m(H') + 5$ as required. \square

Let us now give the algorithm for graphs of girth 5. For a graph G embedded in a surface with boundary cycles C_1, \dots, C_t , let $q(G) = \sum_{i=1}^t \max(7, |C_i|)$.

Lemma 5.5. *For every surface Σ and an integer $k \geq 0$, there exists a quadratic-time algorithm whose input is a triangle-free graph G with closed 2-cell embedding in Σ such that $q(G) \leq k$ and every 4-cycle of G is a boundary cycle, and a 3-coloring ψ of the union B of the boundary cycles of G , and the algorithm decides whether ψ extends to a 3-coloring of G and returns such a 3-coloring if that is the case.*

Proof. We proceed by induction on the complexity of the surface; that is, we assume that such an algorithm exists for all surfaces with smaller genus than Σ , and for the surfaces with the same genus and with fewer cuffs.

Using the algorithm of Theorem 1.3, we can decide whether ψ extends to a 3-coloring of G . Hence, we can assume that this is the case, and we only need to construct a 3-coloring φ of G that extends ψ . If $|V(G)| < R_{\Sigma,k}$, then we can find φ in constant time by brute force; hence, assume that $|V(G)| \geq R_{\Sigma,k}$.

For a surface Π and integer m , let $\nu(\Pi, m) = \frac{5}{s(5)}w_\eta(B) + 18$ (thus, ν is a constant function). We apply the algorithm of Lemma 2.6 to G ; let H be the resulting subgraph of G . Suppose that $H \neq B$. Then note that for every face h of H , the surface Σ_h is at most as complex as Σ . Furthermore, G_h has strictly fewer faces than G , thus ensuring that the recursion eventually terminates. We use the algorithm of Theorem 1.3 to find a 3-coloring ψ' of H such that ψ' extends to a 3-coloring of G , and the restriction of ψ' to B is equal to ψ . By recursive calls to the algorithm we are designing, we find an extension of ψ' to a 3-coloring of G_h for every face h of H , and thus obtain a 3-coloring of G that extends ψ .

Hence, we can assume that $H = B$, and thus G does not contain any connected essential subgraph with fewer than $\nu(\Sigma_h, |V(B)|)$ edges. In particular, if a path P of length at most 4 has both ends in B , then both ends of P belong to the same boundary cycle C . If the distance between the ends of P in C is more than $|P|$, we proceed as in the previous paragraph, with $H = B + P$. Hence, assume that for any such path P , the distance between the ends of P in C is at most $|P|$. Similarly, if G contains a non-contractible cycle K of length at most 7 distinct from the boundary cycles (such a cycle can be found in linear time using the data structure of Lemma 2.3), then K surrounds a cuff, and we can proceed as in the previous paragraph, with $H = B + K$. Hence, assume that no such cycle exists.

Suppose that B contains two adjacent vertices v_1 and v_2 whose degree in G is two. Let C be the boundary cycle containing the edge v_1v_2 . Since the embedding of G is closed 2-cell, C is not a triangle. If Σ is the disk and $|C| \leq 5$, then we can find a 3-coloring of G that extends ψ by the algorithm of Theorem 5.2. If either Σ is not the disk or $|C| \geq 6$, then using the observations of the previous paragraph, we conclude that the edge v_1v_2 is not contained in any (≤ 5) -cycle distinct from C . Let G' be the graph obtained from G by contracting the edge v_1v_2 . Let B' be the boundary cycle of G' and let ψ' be a 3-coloring of B' that matches ψ on $V(B) \setminus \{v_1, v_2\}$. By a recursive call, we obtain a 3-coloring of G' that extends ψ' , and this gives a 3-coloring of G that extends ψ . Hence, assume that B does not contain adjacent vertices whose degree in G is two.

It follows that G satisfies the assumptions of Theorem 5.4. Let us run the algorithm of Theorem 5.4 for G and ψ . Suppose that the algorithm returns a clearable cycle C , bounding an open disk Λ . By a recursive call, we obtain a 3-coloring φ_1 of $G - \Lambda$ that extends ψ . Using algorithm of Theorem 5.2, we find a 3-coloring of G_Λ that extends the restriction of φ_1 to C . By combining these colorings, we obtain a 3-coloring of G that extends ψ .

Finally, suppose that the algorithm returns a reducible configuration F and a reduction G' . By the algorithm of Theorem 4.1 (subdividing the edges in boundary 4-cycles as necessary to accommodate its assumptions), we test whether ψ extends to a 3-coloring of G' . If it does, then by a recursive call we obtain a 3-coloring of G' that extends ψ , and by the properties of reducible configurations, we transform it in constant time in a 3-coloring of G that extends ψ .

Hence, suppose that ψ does not extend to a 3-coloring of G' . In this case,

the algorithm of Theorem 4.1 also returns a subgraph H' of G' such that $B \subseteq H'$, ψ does not extend to a 3-coloring of H' , and $w_\eta(H') \leq w_\eta(B)$. Since H' is triangle-free and all 4-cycles in H' are contained in B , this implies that $|E(H')| \leq \frac{5}{s(5)}w_\eta(B) < \nu(\Sigma, k) - |F|$. It follows that every connected component of H' is non-essential. By repeatedly deleting vertices and edges x of H' not belonging to B such that ψ does not extend to a 3-coloring of $H' - x$, we can assume that ψ extends to a 3-coloring of every proper subgraph of H' that contains B . Since every planar graph of girth at least 5 is 3-colorable, no component of H' is contained in an open disk in Σ . Hence, for each cuff C of Σ , there exists a subset Λ_C of Σ that is an open disk with a hole whose boundary is equal to C , such that for distinct cuffs C and C' , we have $\Lambda_C \cap \Lambda_{C'} = \emptyset$, Λ_C contains a single component H'_C of H' , and each component of H' is contained in one of these sets.

Since $|E(H')| < \nu(\Sigma, k) - |F|$ and F is connected, and since no connected essential subgraph of G has fewer than $\nu(\Sigma, k)$ edges, it follows that at most one component H'_C of H' touches F . Since ψ extends to a 3-coloring of G , it also extends to a 3-coloring of every component of H' which forms a subgraph of G . Since ψ extends to a 3-coloring of every proper subgraph of H' that contains B , but does not extend to a 3-coloring of H' , we conclude that $H' = H'_C \cup B$, and thus H' is near-planar. Also, H' is 2-connected, as otherwise, H'_C contains a 1-cut, and removing the component of the 1-cut that does not contain C from H' results in a proper subgraph that has no 3-coloring extending ψ , which contradicts the choice of H' . The last conclusion of Theorem 5.4 implies that H' is not very exceptional, and by Theorem 5.1, it follows that every face of H'_C in its embedding in Λ_C has length at most $|C| - 5$.

Let H be the subgraph of G with $B \subseteq G$ obtained by the next-to-last conclusion of Theorem 5.4. Note that every face of the main component of H has length at most $|C|$. Using the algorithm of Theorem 1.3, we determine a 3-coloring ψ' of H that extends to a 3-coloring of G and whose restriction to B is equal to ψ . For every face h of H , we determine a 3-coloring of G_h that extends ψ' by a recursive call. By combining these colorings, we obtain a 3-coloring of G that extends ψ .

To analyze the time complexity of the algorithm, note that in each case, we spend a linear time processing G , and we recurse on several graphs G_1, \dots, G_k , such that the sum of the numbers of faces of G_1, \dots, G_k is at most the number of the faces of G . Since the number of faces of G is linear in the number of its vertices, this implies that the time complexity of the algorithm is quadratic. \square

Finally, let us deal with 4-cycles.

Proof of Theorem 1.4. Again, we construct the algorithm by induction on the complexity of the surface. Firstly, we use the algorithm of Theorem 1.3 to determine whether ψ extends to a 3-coloring of G . If it does not, we output this answer and end. Hence, suppose that ψ extends to a 3-coloring of G .

In linear time, we decide whether G contains a non-contractible (≤ 5)-cycle K distinct from the boundary cycles. If that is the case, we let $H = B \cup K$

and we use the algorithm of Theorem 1.3 to find a 3-coloring ψ' of H that extends to a 3-coloring of G and whose restriction to B is equal to ψ . For every face h of H , we determine a 3-coloring of G_h that extends ψ' by a recursive call. By combining these colorings, we obtain a 3-coloring of G that extends ψ . Hence, assume that G contains no non-contractible (≤ 5)-cycle. Similarly, we can assume that every boundary cycle is induced, and that the distance between any two boundary cycles is at least three.

Next, in linear time we decide whether G contains a contractible (≤ 5)-cycle K that does not bound a face. If such a cycle K exists, then let Λ be the open disk bounded by K . We first extend ψ to a 3-coloring of $G - \Lambda$ by a recursive call, and then extend the coloring to G_Λ using the algorithm of Theorem 2.5. Hence, assume that every contractible (≤ 5)-cycle in G bounds a face. Similarly, we can assume that every vertex of G of degree at most two belongs to B . Furthermore, using the ideas of this and the previous paragraph, we can assume that the embedding of G is closed 2-cell.

If G does not contain any 4-face, then we apply the algorithm of Lemma 5.5. Hence, let $f = v_1v_2v_3v_4$ be a 4-face in G . Since the boundary cycles of G are induced and the distance between them is at least three, we can assume that $v_2 \notin V(B)$. Suppose that $v_1, v_3 \in V(B)$. Then both v_1 and v_3 belong to the same boundary cycle C , and since G is triangle-free, the distance between v_1 and v_3 in C is at least two. Hence, if h is a face of $H = B + v_1v_2v_3$, then either Σ_h is strictly less complex than h , or Σ_h is homeomorphic to Σ and all the boundary cycles of G_h are at most as long as the corresponding boundary cycles of G . We use the algorithm of Theorem 1.3 to find a 3-coloring ψ' of H that extends to a 3-coloring of G and whose restriction to B is equal to ψ , and extend ψ' to a 3-coloring of G by recursive calls on the graphs G_h for all faces h of H .

Finally, suppose that $|\{v_1, v_3\} \cap V(B)| \leq 1$. By symmetry, we can assume that $v_1 \notin V(B)$. Using Theorem 1.3, we find a 3-coloring ψ' of $B + v_1v_2v_3v_4$ that extends to a 3-coloring of G and whose restriction to B is equal to ψ . Note that either $\psi'(v_1) = \psi'(v_3)$ or $\psi'(v_2) = \psi'(v_4)$. By symmetry, we can assume the former. Suppose that G contains a path P of length at most 3 joining v_1 with v_3 and disjoint from $\{v_2, v_4\}$, and let K be the cycle $P + v_1v_2v_3$. Since $|K| \leq 5$ and K is distinct from the boundary cycles, it follows that K is contractible, and thus K bounds a face. However, then v_2 has degree two, which is a contradiction. Therefore, there exists no such path, and thus the graph G' obtained from G by identifying v_1 with v_3 is triangle-free. Furthermore, since ψ' extends to a 3-coloring of G , it follows that ψ extends to a 3-coloring of G' . We can find such a 3-coloring of G' by a recursive call, and extend it to G by giving v_1 and v_3 the color of the corresponding vertex of G' .

To analyze the time complexity of the algorithm, note that in each case, we either call the quadratic algorithm of Lemma 5.5, or we spend a linear time processing G and recurse on several graphs G_1, \dots, G_k , such that the sum of the numbers of faces of G_1, \dots, G_k is at most the number of the faces of G . Since the number of faces of G is linear in the number of its vertices, this implies that the time complexity of the algorithm is quadratic. \square

References

- [1] DVOŘÁK, Z., KAWARABAYASHI, K., AND THOMAS, R. Three-coloring triangle-free planar graphs in linear time. *Trans. on Algorithms* 7 (2011), article no. 41.
- [2] DVOŘÁK, Z., KRÁL', D., AND THOMAS, R. Three-coloring triangle-free graphs on surfaces II. 4-critical graphs in a disk. *ArXiv e-prints 1302.2158* (May 2013).
- [3] DVOŘÁK, Z., KRÁL', D., AND THOMAS, R. Three-coloring triangle-free graphs on surfaces III. Graphs of girth five. *ArXiv e-prints 1402.4710* (Feb. 2014).
- [4] DVOŘÁK, Z., KRÁL', D., AND THOMAS, R. Three-coloring triangle-free graphs on surfaces IV. Bounding face sizes of 4-critical graphs. *ArXiv e-prints 1404.6356v2* (Jan. 2015).
- [5] DVOŘÁK, Z., KRÁL', D., AND THOMAS, R. Three-coloring triangle-free graphs on surfaces VI. 3-colorability of quadrangulations. Manuscript, 2015.
- [6] EPPSTEIN, D. Subgraph isomorphism in planar graphs and related problems. *J. Graph Algorithms Appl.* 3 (1999), 1–27.
- [7] GAREY, M., AND JOHNSON, D. *Computers and Intractability: A Guide to the Theory of NP-completeness*. WH Freeman & Co. New York, NY, USA, 1979.
- [8] THOMASSEN, C. Grötzsch's 3-color theorem and its counterparts for the torus and the projective plane. *J. Combin. Theory, Ser. B* 62 (1994), 268–279.
- [9] THOMASSEN, C. The chromatic number of a graph of girth 5 on a fixed surface. *J. Combin. Theory, Ser. B* 87 (2003), 38–71.
- [10] YOUNGS, D. 4-chromatic projective graphs. *Journal of Graph Theory* 21 (1996), 219–227.