Three-coloring triangle-free planar graphs in linear time*

Zdeněk Dvořák† Ken-ichi Kawarabayashi‡

Robin Thomas§

12 August 2008

Abstract

Grötzsch's theorem states that every triangle-free planar graph is 3-colorable, and several relatively simple proofs of this fact were provided by Thomassen and other authors. It is easy to convert these proofs into quadratic-time algorithms to find a 3-coloring, but it is not clear how to find such a coloring in linear time (Kowalik used a nontrivial data structure to construct an $O(n \log n)$ algorithm).

We design a linear-time algorithm to find a 3-coloring of a given triangle-free planar graph. The algorithm avoids using any complex data structures, which makes it easy to implement. As a by-product we give another simple proof of Grötzsch's theorem.

1 Introduction

The following is a classical theorem of Grötzsch [6].

Theorem 1.1. *Every triangle-free planar graph is 3-colorable.*

*An extended abstract of this paper was submitted to SODA'09.

†Department of Applied Mathematics, Charles University, Prague, Czech Republic.

‡National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan.

§School of Mathematics, Georgia Institute of Technology, Atlanta, Georgia 30332-0160, USA. Partially supported by NSF under Grant No. DMS-0701077.

This result has been the subject of extensive research. Thomassen [14, 15] found two short proofs and extended the result in many ways. We return to the various extensions later, but let us discuss algorithmic aspects of Theorem 1.1 first. It is easy to convert either of Thomassen's proofs into a quadratic-time algorithm to find a 3-coloring, but it is not clear how to do so in linear time. A serious problem appears very early in the algorithm. Given a facial cycle C of length four, one would like to identify a pair of diagonally opposite vertices of C and apply recursion to the smaller graph. It is easy to see that at least one pair of diagonally opposite vertices on C can be identified without creating a triangle, but how can we efficiently decide which pair? If we could test in (amortized) constant time whether given two vertices are joined by a path of length at most three, then that would take care of this issue. This can, in fact, be done, using a data structure of Kowalik and Kurowski [8] *provided* the graph does not change. In our application, however, we need to repeatedly identify vertices, and it is not clear how to maintain the data structure of Kowalik and Kurowski in overall linear time. Kowalik [7] developed a sophisticated enhancement of this data structure that supported vertex identification, but at the expense of an added $\log n$ factor. Thus he designed an $O(n \log n)$ algorithm to 3-color a triangle-free planar graph on n vertices. We improve this to a linear-time algorithm, as follows.

Theorem 1.2. *There is a linear-time algorithm to* 3*-color an input triangle-free planar graph.*

To describe the algorithm we exhibit a specific list of five reducible configurations, called "multigrams", and show that every triangle-free planar graph contains one of those reducible configurations. Proving this is the only step that requires some effort; the rest of the algorithm is entirely straightforward, and the algorithm is very easy to implement. Given a triangle-free planar graph G we look for one of the reducible configurations in G, and upon finding one we modify G to a smaller graph G', and apply the algorithm recursively to G'. It is easy to see that every 3-coloring of G' can be converted to a 3-coloring of G in constant time. Furthermore, each reducible configuration has a vertex of degree at most three, and, conversely, given a vertex of G of degree at most three it can be checked in constant time whether it belongs to a reducible configuration. Thus at every step a reducible configuration can be found in amortized constant time by maintaining a list of candidates for such vertices. As a by-product of the proof of correctness of our algorithm we give a short proof of Theorem 1.1.

Let us briefly survey some of the related work. Since in a proof of Theorem 1.1 it is easy to eliminate faces of length four, the heart of the argument lies in proving the theorem for graphs of girth at least five. For such graphs there are several extensions of the theorem. Thomassen proved in [14] that every graph of girth at least five that admits an embedding in the projective plane or the torus is 3-colorable, and the analogous result for Klein bottle graphs was obtained in [13]. For a general surface Σ Thomassen [16] proved the deep theorem that there are only finitely many 4-critical graphs of girth at least five that embed in Σ. (A graph is 4-critical if it is not 3-colorable, but every proper subgraph is.)

None of the results mentioned in the previous paragraph hold without the additional restriction on girth. Nevertheless, Gimbel and Thomassen [5] found an elegant characterization of 3-colorability of triangle-free projective-planar graphs. That result does not seem to extend to other surfaces, but two of us in joint work with Kráľ [3] were able to find a sufficient condition for 3-colorability of triangle-free graphs drawn on a fixed surface Σ. The condition is closely related to the sufficient condition for the existence of disjoint connecting trees in [11]. Using that condition Dvořák, Kráľ and Thomas were able to design a linear-time algorithm to test if a triangle-free graph on a fixed surface is 3-colorable [3].

If we allow the planar graph G to have triangles, then testing 3-colorability becomes NP-hard [4]. There is an interesting conjecture of Steinberg stating that every planar graph with no cycles of length four or five is 3-colorable, but that is still open. Every planar graph is 4-colorable by the Four-Color Theorem [1, 2, 10], and a 4-coloring can be found in quadratic time [10]. Any improvement to the running time of this algorithm would seem to require new ideas. A 5-coloring of a planar graph can be found in linear time [9].

Our terminology is standard. All *graphs* in this paper are simple and *paths* and *cycles* have no repeated vertices. By a *plane graph* we mean a graph that is drawn in the plane. On several occasions we will be identifying vertices, but when we do, we will remove the resulting parallel edges. When this will be done by an algorithm we will make sure that the only parallel edges that arise will form faces of length two. The detection and removal of such parallel edges can be done in constant time.

2 Short proof of Grötzsch's theorem

Let G be a plane graph. By a *tetragram* in G we mean a sequence (v_1, v_2, v_3, v_4) of vertices of G such that they form a facial cycle in G in the order listed. We define a *hexagram* $(v_1, v_2, \ldots, v_6)$ similarly. By a *pentagram* in G we mean a sequence $(v_1, v_2, v_3, v_4, v_5)$ of vertices of G such that they form a facial cycle in G in the order listed and v_1, v_2, v_3, v_4 all have degree exactly three. We will show that every triangle-free planar graph of minimum degree at least three has a tetra-, penta- or hexagram with certain additional properties that will allow an inductive argument. But first we need the following lemma.

Lemma 2.1. *Let G be a connected triangle-free plane graph, let C be the facial cycle in G bounding the unbounded face f_0 of G, assume that C has length at most six, and assume that every vertex of G not on C has degree in G at least three. If $G \neq C$, then G has either a tetragram, or a pentagram $(v_1, v_2, v_3, v_4, v_5)$ such that $v_1, v_2, v_3, v_4 \notin V(C)$.*

Proof. We may assume that G has no separating cycle C' of length at most six, for otherwise we can apply the lemma to C' and the subgraph of G consisting of all vertices and edges drawn in the closed disk bounded by C'. We define the charge of a vertex v to be $3\deg(v) - 12$, the charge of the face f_0 to be $3|V(C)| + 11$ and the charge of a face $f \neq f_0$ of length l to be $3l - 12$. It follows from Euler's formula that the sum of the charges of all vertices and faces is -1.

We now redistribute the charges according to the following rules. Every vertex not on C of degree three will receive one unit of charge from each incident face, each vertex on C of degree three will receive three units from f_0, and each vertex of degree two on C will receive five units from f_0 and one unit from the other incident face. Thus the final charge of every vertex is non-negative.

We now show that the final charge of f_0 is also non-negative. Let l denote the length of C. Then f_0 has initial charge of $3l + 11$. By hypothesis at least one vertex of C has degree at least three, and hence f_0 sends a total of at most $5(l - 1) + 3$ units of charge, leaving it at the end with charge of at least $3l + 11 - 5(l - 1) - 3 \geq 1$.

Since no charge is lost or created, there is a face $f \neq f_0$ whose final charge is negative. Since f sends at most one unit to each incident vertex, we see that f has length at most five. Furthermore, if f has length exactly five, then it sends one unit to at least four incident vertices. None of those could be a degree two vertex on C, for then f would not be sending anything to the ends of the common subpath of the boundaries of f and f_0. Thus the vertices of f form the desired tetragram or pentagram. $\square$

Let $k = 4, 5, 6$, and let $(v_1, v_2, \ldots, v_k)$ be a tetragram, pentagram or hexagram in a triangle-free plane graph G. If $k = 4$ or $k = 6$, then we say that $(v_1, v_2, \ldots, v_k)$ is *safe* if every path in G of length at most three with ends v_1 and v_3 is a subgraph of the cycle $v_1 v_2 \cdots v_k$. For $k = 5$ we define safety as follows. For $i = 1, 2, 3, 4$ let x_i be the third neighbor of v_i. Then $x_i \notin V(C)$, because G is triangle-free. Assume that the vertices x_1, x_2, x_3, x_4 are pairwise distinct and pairwise non-adjacent. Let either $i = 2$ and $j = 1$, or $i = 3$ and $j = 4$. Assume that there is no path in $G \backslash \{v_1, v_2, v_3, v_4\}$ of length at most three from x_i to v_5, and assume that every path in $G \backslash \{v_1, v_2, v_3, v_4\}$ of length at most three from x_i to x_j has length exactly two, and its completion via the path $x_j v_j v_i x_i$ results in a facial cycle of G. In those circumstances we say that the pentagram $(v_1, v_2, \ldots, v_5)$ is *safe*.

Lemma 2.2. *Every triangle-free plane graph of minimum degree at least three has a safe tetragram, a safe pentagram, or a safe hexagram.*

Proof. Let G be as stated. If G has a separating cycle of length at most six, then let us select such a cycle C so that the disk it bounds as small as possible, and let G' be the subgraph of G consisting of all vertices and edges drawn in the closed disk bounded by C. If G has no separating cycle of length at most six, then let $G' := G$ and let C be an arbitrary facial cycle in G of length at most five. It is easy to see that if (v_1, v_2, v_3, v_4) is a tetragram in G, then one of the tetragrams (v_1, v_2, v_3, v_4), (v_2, v_3, v_4, v_1) is safe. Thus we may assume that G has no tetragram. ¿From Lemma 2.1 applied to the graph G' and facial cycle C we deduce that G' has a pentagram $(v_1, v_2, v_3, v_4, v_5)$ such that $v_1, v_2, v_3, v_4 \notin V(C)$. Given the choice of C and the fact that G has no tetragram, it follows that this pentagram safe, unless (up to symmetry) there is a path $x_3 a b x_4$ in $G \backslash \{v_1, v_2, v_3, v_4\}$ forming a facial hexagon $x_4 v_4 v_3 x_3 a b$. But then the choice of C implies that $(x_4, v_4, v_3, x_3, a, b)$ is a safe hexagram, as desired. $\square$

Proof of Theorem 1.1. Let G be a triangle-free plane graph. We proceed by induction on $|V(G)|$. We may assume that every vertex v of G has degree at least three, for otherwise the theorem follows by induction applied to $G \backslash v$. By Lemma 2.2 there is a safe tetra-, penta-, or hexagram $(v_1, v_2, \ldots, v_k)$.

If $k = 4$ or $k = 6$, then we apply induction to the graph obtained from G by identifying v_1 and v_3. It follows from the definition of safety that the new graph has no triangles, and clearly every 3-coloring of the new graph extends to a 3-coloring of G. Thus we may assume that $(v_1, v_2, \ldots, v_5)$ is a safe pentagram in G. Let G' be obtained from $G \setminus \{v_1, v_2, v_3, v_4\}$ by identifying v_5 with x_2, and x_3 with x_4. It follows from the definition of safety that G' is triangle-free, and hence it is 3-colorable by the induction hypothesis. It is routine to verify that any 3-coloring of G' can be extended to a 3-coloring of G. Thus G is 3-colorable, as desired. □

Let us note that the essential ideas of the proof came from Thomassen's work [14]. For graphs of girth at least five Thomassen actually proves a stronger statement, namely that every 3-coloring of an induced facial cycle of length at most nine extends to a 3-coloring of the entire triangle-free plane graph, unless some vertex of G has three distinct neighbors on C (and those neighbors received three different colors). By restricting ourselves to Theorem 1.1 we were able to somewhat streamline the argument. Another variation of the same technique is presented in [7].

3 Graph representation

In the algorithm, we represent the planar graph in the following (standard) way: Each edge is represented by a pair of opposite directed edges. The directed edges are partitioned into closed directed walks, forming the boundaries of the faces. All these walks are oriented in a consistent way, so that if a directed edge e belongs to the directed walk bounding a face f, then f is on the left side of e.

For each directed edge e, we maintain pointers to the corresponding opposite edge and to the preceding and following edge of the directed boundary walk to that e belongs, as well as the two vertices it joins. For each vertex, we record a pointer to one outgoing edge. It is not necessary to maintain lists of incoming and outgoing edges, as these are represented implicitly; e.g., to list the outgoing edges incident with a vertex v, we take the single outgoing edge e_1 that we have recorded, find the opposite edge e_1', take the edge e_2 following e_1' in the directed facial walk, and repeat this process until we obtain all the edges e_1, e_2, ... incident with v.

Suppose that D is a fixed constant (in our algorithm, $D = 47$). We can perform the following operations with graphs represented in the described way in constant time:

- remove an edge
- add an edge, assuming that the edges preceding and following it in the facial walks are specified
- remove an isolated vertex
- determine the degree of a vertex v if $\deg(v) \leq D$, or prove that $\deg(v) > D$
- check whether two vertices u and v such that $\min(\deg(u), \deg(v)) \leq D$ are adjacent
- check whether the distance between two vertices u and v such that $\max(\deg(u), \deg(v)) \leq D$ is at most two
- determine the length ℓ of a face f incident with an edge e if $\ell \leq D$, or prove that $\ell > D$
- list a subgraph consisting of vertices reachable from a vertex v_0 through a path v_0, v_1, $\ldots$, v_t of length $t \leq D$, such that $\deg(v_i) \leq D$ for $0 \leq i < t$ (but the degree of v_t may be arbitrary).

All the transformations and queries executed in the algorithm can be expressed in terms of these simple operations.

4 The algorithm

The idea of our algorithm is to find a safe tetragram, pentagram or hexagram γ in G and use it to reduce the size of the graph as in the proof of Theorem 1.1 above. Finding γ is easy, but the difficulty lies in testing safety. To resolve this problem we prove a variant of Lemma 2.2 that will guarantee the existence of such γ with an additional property that will allow testing safety in constant time. The additional property, called security, is merely that enough vertices in and around γ have bounded degree. Unfortunately, the additional property we require necessitates the introduction of another configuration, a variation of pentagram, called "decagram". For the sake of consistency, we say that a *monogram* in a graph G is the one-vertex sequence (v) comprised of a vertex $v \in V(G)$ of degree at most two.

Now let G be a plane graph, let $k \in \{1, 4, 5, 6\}$ and let $\gamma = (v_1, v_2, \ldots, v_k)$ be a mono-, tetra-, penta-, or hexagram in G. Let C be a subgraph of G.

(For the purpose of this section the reader may assume that C is the null graph, but in the next section we will need C to be a facial cycle of G.) A vertex of G is *big* if it has degree at least 48, and *small* otherwise. A vertex $v \in V(G)$ is *C-admissible* if it is small and does not belong to C; otherwise it is *C-forbidden*. A pentagram $(v_1, v_2, \ldots, v_5)$ is called a *decagram* if v_5 has degree exactly three (and hence $v_1, \ldots, v_5$ all have degree three). A *multigram* is a monogram, tetragram, pentagram, hexagram or a decagram. The vertex v_1 will be called the *pivot* of the multigram $(v_1, v_2, \ldots, v_k)$. In the following γ will be a multigram, and we will define (or recall) what it means for γ to be safe and C-secure. We will also define a smaller graph G', which will be called the *γ-reduction of G*.

If γ is a mongram, then we define it to be always *safe*, and we say that it is *C-secure* if $v_1 \notin V(C)$. We define $G' := G \backslash v_1$.

Now let γ be a tetragram. Let us recall that γ is safe if the only path in G of length at most three with ends v_1 and v_3 is a subgraph of the facial cycle $v_1v_2v_3v_4$. We say that γ is *C-secure* if it is safe, v_1 is C-admissible and has degree exactly three, and letting x denote the neighbor of v_1 other than v_2 and v_4, the vertex x is C-admissible, and either v_3 is C-admissible, or every neighbor of x is C-admissible. We define G' to be the graph obtained from G by identifying the vertices v_1 and v_3 and deleting one edge from each of the two pairs of parallel edges that result.

Now let γ be a decagram, and for $i = 1, 2, 3, 4$ let x_i be the neighbor of v_i other than v_{i-1} or v_{i+1}, where v_0 means v_5. We say that the decagram γ is *safe* if x_1, x_3 are distinct, non-adjacent and there is no path of length two between them. We say that γ is *C-secure* if it is safe and the vertices $v_1, v_2, \ldots, v_5, x_1, x_3$ are all C-admissible. We define G' to be the graph obtained from $G \backslash \{v_1, v_2, \ldots, v_5\}$ by adding the edge x_1x_3.

Now let γ be a pentagram, and for $i = 1, 2, 3, 4$ let x_i be as in the previous paragraph. Let us recall that the safety of γ was defined prior to Lemma 2.2. We say that γ is *C-secure* if it is safe, the vertices $v_1, v_2, \ldots, v_5, x_1, x_2, x_3, x_4$ are all C-admissible, either v_5 or x_2 has no C-forbidden neighbor, and either x_3 or x_4 has no C-forbidden neighbor. We define G' as in the proof of Theorem 1.1: G' is obtained from $G \backslash \{v_1, v_2, v_3, v_4\}$ by identifying x_2 and v_5; identifying x_3 and x_4; and deleting one of the parallel edges should x_3 and x_4 have a common neighbor.

Finally, let γ be a hexagram. Let us recall that γ is safe if every path of length at most three in G between v_1 and v_3 is the path $v_1v_2v_3$. We say that γ is *C-secure* if v_1, v_3, v_6 are C-admissible, v_1 has degree exactly three, and the neighbor of v_1 other than v_2 or v_6 is C-admissible. We define G'

to be the graph obtained from G be identifying the vertices v_1 and v_3 and deleting one of the parallel edges that result.

We say that a multigram γ is *secure* if it is K_0-secure, where K_0 denotes the null graph. This completes the definition of safe and secure multigrams.

Lemma 4.1. *Let G be a triangle-free plane graph, let γ be a safe multigram in G, and let G' be the γ-reduction of G. Then G' is triangle-free, and every 3-coloring of G' can be converted to a 3-coloring of G in constant time. Moreover, if γ is secure, then G' can be regarded as having been obtained from G by deleting at most* 100 *edges, adding at most* 91 *edges, and deleting at least one isolated vertex.*

Proof. The graph G' is triangle-free, because γ is safe. A simple case-checking shows that every 3-coloring of G' can be extended to a 3-coloring of G. If γ is secure, then every time vertices u and v are identified in the construction of G', one of u, v is small. Thus the identification of u and v can be seen as a deletion of at most 47 edges and addition of at most 47 edges. The theorem follows by a more careful examination of the construction of G'. □

Let G and C be as above. We say that two vertices $u, v \in V(G)$ are *close* if there exists a path P in G with ends u and v of length at most four such that every vertex of P (including its ends) is small. Thus for every vertex v there are at most $1 + 47 + 47^2 + 47^3 + 47^4$ vertices that are close to v.

Lemma 4.2. *Given a triangle-free plane graph G and a vertex $v \in V(G)$, it can be decided in constant time whether G has a secure multigram with pivot v.*

Proof. This follows by inspecting the subgraph of G induced by vertices that are close to v. To test safety we need to check the existence of certain paths P of bounded length with prescribed ends. However, whenever such a test is needed every vertex of P, except possibly one, is small. Thus the test can be carried out in constant time. □

Lemma 4.3. *Let G and G' be triangle-free plane graphs, such that for some pair of non-adjacent vertices $u, v \in V(G)$ the graph G' is obtained from G by adding the edge uv. Let γ be a secure multigram in exactly one of the graphs G, G'. Then the pivot of γ is close to u or v in G.*

Proof. This follows from the fact that all vertices that impact the security of a multigram are close to its pivot. □

The next theorem will serve as the basis for the proof of correctness of our algorithm. We defer its proof until the next section.

Theorem 4.4. *Every non-null triangle-free planar graph has a secure multigram.*

We are now ready to prove Theorem 1.2, assuming Theorem 4.4.

Algorithm 4.5. *There is an algorithm with the following specifications:*
Input: *A triangle-free planar graph.*
Output: *A proper* 3*-coloring of* G*.*
Running time: $O(|V(G)|)$.

Description. Using a linear-time planarity algorithm that actually outputs an embedding, such as [12] or [17], we can assume that G is a plane graph. The algorithm is recursive. Throughout the execution of the algorithm we will maintain a list L that will include the pivots of all secure multigrams in G, and possibly other vertices as well. We initialize the list L to consist of all vertices of G of degree at most three.

At a general step of the algorithm we remove a vertex v from L. There is such a vertex by Theorem 4.4 and the requirement that L include the pivots of all secure multigrams. If $v \notin V(G)$, then we go to the next iteration. Otherwise, we check if G has a secure multigram with pivot v. This can be performed in constant time by Lemma 4.2. If no such multigram exists, then we go to the next iteration. Otherwise, we let γ be one such multigram, and let G' be the γ-reduction of G. By Lemma 4.1 the graph G' is triangle-free and can be constructed in constant time by adding and deleting bounded number of edges. For every edge uv that was deleted or added during the construction of G' we add to L all vertices that are close to u or v. By Lemma 4.3 this will guarantee that L will include pivots of all secure multigrams in G'. We apply the algorithm recursively to G', and convert the resulting 3-coloring of G' to one of G using Lemma 4.1. Since the number of vertices added to L is proportional to the number of vertices removed from G we deduce that the number of vertices added to L (counting multiplicity) is at most linear in the number of vertices of G. Thus the running time is $O(|V(G)|)$, as claimed.

Algorithm 4.5 has the following extension.

Algorithm 4.6. *There is an algorithm with the following specifications:*
Input: *A triangle-free plane graph G, a facial cycle C in G of length at most five, and a proper 3-coloring ϕ of C.*
Output: *A proper 3-coloring of G whose restriction to $V(C)$ is equal to ϕ.*
Running time: $O(|V(G)|)$.

Description. The description is exactly the same, except that we replace "secure" by "C-secure" and appeal to Lemma 5.1 rather than Theorem 4.4.

5 Proof of correctness

In this section we prove Theorem 4.4, thereby completing the proof of correctness of the algorithm from the previous section. The theorem will follow from the next lemma. Unfortunately, for a technical reason we need a small variation on the notion of C-secure tetragram. Let G be a triangle-free plane graph, let C be a cycle in G, let H be the subgraph of G consisting of all vertices and edges of G drawn in the closed disk bounded by C, and let $\gamma = (v_1, v_2, v_3, v_4)$ be a safe tetragram in H. We say that γ is a *C-semigram* if v_1 is C-admissible, has degree exactly three and belongs to H, the neighbor of v_1 other than v_2 and v_4 is also C-admissible, and the edge v_3v_4 belongs to C. We say that v_3 is the *tail of the C-semigram γ*. We say that a vertex $v \in V(G)$ is a *C-appendix* if either $v \in V(H) - V(C)$ and v is big, or $v \in V(C)$ and v is the tail of some K-semigram in H for some cycle K in H. If xy is an edge in a plane graph, and f is a face of G incident with y but not with the edge xy, then we say that *f is opposite to xy*. Let us emphasize that this notion is not symmetric in x, y.

Lemma 5.1. *Let G be a connected triangle-free plane graph, let C be the facial cycle bounding the outer face f_0, and assume that C has at most six vertices and that $|V(G) - V(C)| \geq 2$. Then either G has a C-secure multigram, or C has length exactly six and includes at least two distinct non-adjacent C-appendices.*

Proof. Suppose for a contradiction that the lemma is false, and let G be a counterexample with $|V(G)|$ minimum. We first establish the following claim.

(1) *If K is an induced separating cycle in G of length at most six such that the open disk bounded by K includes at least two vertices of G,*

then K has length exactly six and includes two distinct non-adjacent C-appendices.

To prove (1) let K be as stated, and let G' be the subgraph of G consisting of all vertices and edges that belong to the closed disk bounded by K. ¿From the induction hypothesis applied to G' and K we deduce that either G' has a K-secure multigram, or K has length exactly six and includes at least two different non-adjacent K-appendices u_1, u_2. Since every K-secure multigram in G' is a C-secure multigram in G, we may assume the latter. Now if $u_i \notin V(C)$, then u_i is big, for otherwise the semigram with u_i as tail is a C-secure tetragram (because K is induced), contrary to the fact that G is a counterexample to the theorem. Thus u_i is a C-appendix, and the same conclusion follows if $u_i \in V(C)$. This proves (1).

It follows from (1) that every cycle of length at most five bounds a face, and that C is an induced cycle. It also follows that every tetragram is safe.

We assign charges to vertices and faces of G as follows. Initially, a vertex v will receive a charge of $9\deg(v) - 36$ if $v \notin V(C)$, and $8\deg(v) - 19$ otherwise. The outer face f_0 will receive a charge of zero, and every other face f of length l will receive a charge of $9l - 36$. By Euler's formula the sum of the charges is equal to

$$\sum_{v \notin V(C)} 9(\deg(v) - 4) + \sum_{v \in V(C)} (8\deg(v) - 19) + \sum_{f \neq f_0} 9(\,\mathrm{size}(f) - 4)$$

$$= \sum_{v \in V(G)} 9(\deg(v) - 4) + \sum_{f} 9(\mathrm{size}(f) - 4) - \sum_{v \in V(C)} \deg(v) + 8|V(C)| + 36$$

$$= 8|V(C)| - \sum_{v \in V(C)} \deg(v) - 36 \leq -1,$$

because all vertices of C have degree at least two, and at least one has degree at least three by hypothesis. Furthermore,

(2) *if at least three vertices of C have degree at least three, then the sum of the charges is at most -3.*

We now redistribute the charges according to the following rules. The new charge thus obtained will be referred to as the *final* charge. We need a definition first. Let $f \neq f_0$ be a face of G incident with a vertex $v \in V(C)$. If there exist two consecutive edges in the boundary of f such that both are incident with v and neither belongs to C, then we say that f is a *v-interior face*. The rules are:

(A) every face other than f_0 sends three units of charge to every incident vertex v such that either $v \in V(C)$ and v has degree two in G, or $v \notin V(C)$ and v has degree exactly three,

(B) every big vertex not on C sends three units to each incident face and for every C-semigram (v_1, v_2, v_3, v_4) it sends three units to the face bounded by v_1, v_2, v_3, v_4,

(C) every vertex $v \in V(C)$ sends three units to every v-interior face,

(D) if $x \in V(G)$ is C-forbidden, and y is a C-admissible neighbor of x of degree three, then x sends three units to the unique face opposite to xy, and one unit to the face opposite to yz for every C-admissible neighbor z of y of degree three,

(E) every C-forbidden vertex sends five units to every C-admissible neighbor of degree at least four,

(F) for every C-admissible vertex y of degree at least four that has a C-forbidden neighbor we select a C-forbidden neighbor x of y and send one unit to each face opposite to xy, and one unit to the face opposite to yz for every C-admissible neighbor z of y of degree three.

Since G does not satisfy the conclusion of the theorem, it follows that every vertex of G has degree at least two, every vertex of degree exactly two belongs to C, and there are at most four C-semigrams in G. With these facts in mind we now show that every vertex has non-negative charge. To that end let $v \in V(G)$ have degree d, and assume first that v is C-admissible. If $d = 3$, then it starts out with a charge of -9 and receives three from each incident face by rule (A) for a final total of zero. If $d \geq 4$, then v starts out with a charge of $9d - 36 \geq 0$. If v has no C-forbidden neighbor, then it sends no charge and the claim holds. Thus we may assume that v has a C-forbidden neighbor, and let x be such neighbor selected by rule (F). Then v receives at least five units by rule (E), and sends at most $2d - 3$ by rule (F) for a total of at least $9d - 36 + 5 - (2d - 3) = 7d - 28 \geq 0$. Thus every C-admissible vertex has non-negative final charge. If v is big, but does not belong to C, then it sends only by rules (B), (D) or (E). It sends at most $3d$ using the first clause of rule (B) and at most 12 using the second clause (because there are at most four C-semigrams), and it sends at most $5d$ using rules (D) or (E) for a total final charge of at least $9d - 36 - 3d - 12 - 5d \geq 0$, because $d \geq 48$. Thus we may assume that

$v \in V(C)$. Then v starts out with $8d-19$ and sends $3(d-3)$ using rule (A) or (C) (if $d=2$, then v receives 3 by rule (A); and otherwise it sends $3(d-3)$ by rule (C)) and it sends $5(d-2)$ using rule (D) or (E) for a total of $8d-19-3(d-3)-5(d-2)=0$. This proves our claim that the final charge of every vertex is non-negative.

It also follows that every face of length $l \geq 6$ has non-negative final charge, for every face sends at most three units to each incident vertex and only to those vertices; thus the final charge is at most $9l-36-3l \geq 0$.

We have thus shown that G has a face f of length at most five with strictly negative final charge. Clearly f is not the outer face.

(3) *No vertex incident with f has degree two.*

To prove (3) suppose for a contradiction that a vertex v of degree two is incident with f. Thus v and the two edges incident with v and f belong to C. Since $G \neq C$ and f has length at most five we deduce that at least two vertices incident with f are incident with C and have degree at least three. Those two vertices do not receive any charge from f, and hence f has length four, because it has negative charge.

We deduce that f is bounded by a cycle $u_1u_2u_3u_4$, where u_1, u_2, u_3 are consecutive vertices of C, and u_2 has degree two. It follows that $u_4 \notin V(C)$, because C is induced. Since f has negative charge we find that u_4 is small, and hence C-admissible. Thus the cycle C' obtained from C by replacing the vertex u_2 by u_4 either has length at most five, or does not have two distinct non-adjacent C-appendices. It follows from (1) that the open disk bounded by C' includes at most one vertex, and hence it contains exactly one, because $|V(G)|-V(C)| \geq 2$. Let that vertex be v_4; then the remaining vertices of C can be numbered v_1, v_2, v_3 so that the cycle C is $u_1u_2u_3v_1v_2v_3$ and v_4 is adjacent to v_1, v_3, v_4. Then (v_4, v_1, v_2, v_3) and (u_4, u_1, u_2, u_3) are C-semigrams with distinct non-adjacent tails, contrary to the assumption that G is a counterexample to the theorem. This proves (3).

Let uv be an edge of G such that f is opposite to uv. Let us say that v is a *sink* if v has degree three and both u and v are C-admissible. Let us say that v is a *source* if either $v \notin V(C)$ and v is big, or $v \in V(C)$ and f is v-interior. Since v does not have degree two by (3) we deduce that v is a sink if and only if it receives three units of charge from f by rule (A) and f does not receive three units by rule (D) from the unique neighbor of v opposite to f. Likewise, the vertex v is a source if and only if it sends three units to f by the first clause of rule (B) or by rule (C). Let s be the number of sources, and t the number of sinks. Thus the charge of f is at least $9+3s-3t$ if f has length five and at least $3s-3t$ if f has length four.

Let us assume now that f has length five, and let $v_1, v_2, \dots, v_5$ be the incident vertices, listed in order. Since f has negative charge, at least four of the five incident vertices are sinks, and so we may assume that v_1, v_2, v_3, v_4 are sinks. Thus $\gamma = (v_1, v_2, \dots, v_5)$ is a pentagram. For $i = 1,2,3,4$ let x_i be the third neighbor of v_i. From (1) and the fact that G has no C-secure tetragram we deduce that the vertices x_1, x_2, x_3, x_4 are distinct and pairwise non-adjacent. If v_5 is a sink as well, then it follows from (1) that γ is C-secure decagram. (If there is a path of length two between x_1 and x_3, then G has a separating cycle K of length six such that all vertices of K except possibly one are C-admissible. Thus the open disk bounded by K includes exactly one vertex of G, necessarily of degree three. That vertex is x_2 and it follows that x_2 is adjacent to x_1 and x_3, a contradiction.) Thus v_5 is not a sink, and hence the final charge of f is at least -3. It follows that v_5 is not a source, which in turn implies that v_5 is C-admissible (because v_1 and v_4 are C-admissible). We claim that γ is a safe pentagram. If there exists a path P in $G \backslash \{v_1, v_2, v_3, v_4\}$ of length at most three with ends x_2 and v_5, then P can be completed to a separating cycle K using the path $v_5 v_1 v_2 x_2$. By (1) this cycle bounds an open disk that contains only the vertex x_1, which is impossible, because x_1 is not adjacent to x_2. In order to complete the proof that γ is safe it suffices to consider, by symmetry, a path in $G \backslash \{v_1, v_2, v_3, v_4\}$ of length at most three with ends x_1 and x_2. This path can be completed via the path $x_2 v_2 v_1 x_1$ to a cycle K'. Since v_1 and v_2 have degree three, and x_1 is not adjacent to x_2, we deduce that K' is a facial cycle. Since x_1 is not adjacent to x_2 we may assume that K' has length six; let its vertices in order by $x_1 v_1 v_2 x_2 ab$. Then $(v_1, v_2, x_2, a, b, x_1)$ is a C-secure hexagram in G, a contradiction. This proves our claim that γ is a safe pentagram. We have already established that the vertices $v_1, v_2, \dots, v_5, x_1, x_2, x_3, x_4$ are C-admissible. If x_i has a C-forbidden neighbor for some $i \in \{1,2,3,4\}$, then f receives one unit of charge either from that neighbor by rule (D) if x_i has degree three, or from x_i by rule (F) otherwise. If v_5 has a C-forbidden neighbor, then it sends one unit of charge to f by rule (F). Thus at most two vertices among v_5, x_1, x_2, x_3, x_4 have a C-forbidden neighbor, and hence it follows that either γ, or $(v_4, v_3, v_2, v_1, v_5)$ is a C-secure pentagram, a contradiction.

Thus we have shown that f has length four. Let v_1, v_2, v_3, v_4 be the incident vertices listed in order. Since f has negative charge at least $3s-3t$, we may assume that v_1 is a sink and v_3 is not a source. Since v_3 is not a source and (v_1, v_2, v_3, v_4) is not a C-secure tetragram, (3) implies that exactly one of v_2v_3, v_3v_4 is an edge of C, and hence we may assume the latter. In particular, $v_2 \notin V(C)$. Since we have shown that every tetragram is safe it follows that $\gamma = (v_1, v_2, v_3, v_4)$ is a C-semigram. It is possible that (v_2, v_1, v_4, v_3) is also a C-semigram, in which case f has charge at least -6; otherwise it has charge at least -3. Thus we have shown that

(4) *the charge of every face is at least -3 times the number of C-semigrams its incident vertices give rise to.*

It follows that there are no big vertices, for every big vertex sends three units per C-semigram to the corresponding face by the second clause of rule (B).

Let v be the neighbor of v_1 other than v_2 and v_4. If v has no C-forbidden neighbor, then γ is a C-secure tetragram, a contradiction. Thus v has a neighbor $u \in V(C)$. This has two implications. First, v sends one unit to f by rule (D), which was not accounted for in (4), and hence the constant -3 in (4) is improved to -2. Second, since $v_3, v_4, u \in V(C)$ have degree at least three, the total charge is at most -3 by (2). We deduce that G has at least two C-semigrams.

Let us go back to the C-semigram γ for a little while. Let C, C_1, C_2 be the three cycles in the graph consisting of C and the path uvv_1v_4, numbered so that v_3 belongs to C_2. We claim that C_2 has length at least seven. To prove this claim we note that by (1) the cycle C_2 has length at least six. Assume that C_2 has length exactly six. Then the open disk it bounds contains v_2 and no other vertex of G. It follows that v_2 has degree three, and its third neighbor is u. Since v has degree at least three, it is incident with an edge belonging to the open disk bounded by C_1. If such an edge is a chord, say vw, then we have determined G completely, and the C-semigram (v_1, v, w, v_4) shows that w is a C-appendix, a contradiction. Thus the cycle C_1 is also separating. By (1) the open disk it bounds includes exactly one vertex, a degree three neighbor of v. This determines G completely, and we see that the two neighbors of v_4 on C are distinct non-adjacent C-appendices, a contradiction. This proves our claim that C_2 has length at least seven. Thus C_1 has length at most five, and hence it bounds a face. Since u is not a C-appendix (because there are no two distinct non-adjacent C-appendices) it follows that C_1 has length exactly five. Thus u and v_4 have a common neighbor of degree two on C, say z. Let $f(\gamma)$ denote the face bounded by C_1.

The face $f(\gamma)$ starts out with a charge of 9, sends three units to each of v_1, v, z by rule (A), and receives one either from v_3 by rule (D), or from v_2 by rule (F) for a total of +1. We will use this +1 to partially offset the negative charge of f. To do that we notice that since every two distinct C-appendices are adjacent we have $f(\gamma) \neq f(\gamma')$ for every two distinct C-semigrams γ, γ'. It follows that there are at most two distinct C-semigrams. On the other hand, the C-semigram γ contributes -2 toward f and $+1$ toward $f(\gamma)$, for a grand total of -1, and the same applies to every other C-semigram. Since the total charge is at most -3, there are at least three C-semigrams, a contradiction. □

Proof of Theorem 4.4. Let G be a triangle-free planar graph. We may assume that G is actually drawn in the plane. If G has a vertex of degree two or less, then it has a secure monogram, and so we may assume that G has minimum degree at least three. It follows that G has a facial cycle C of length at most five. Let H be the component of G containing C. We may assume that C bounds the outer face of H. Since H has minimum degree at least three and is triangle-free it follows that $V(H) - V(C)$ has at least two vertices. By Lemma 5.1 H has a C-secure multigram; but any C-secure multigram in H is a secure multigram in G, as desired. □

References

[1] K. Appel and W. Haken, Every planar map is four colorable, Part I: discharging, *Illinois J. of Math.* **21** (1977), 429–490.

[2] K. Appel, W. Haken and J. Koch, Every planar map is four colorable, Part II: reducibility, *Illinois J. of Math.* **21** (1977), 491–567.

[3] Z. Dvořák, D. Král' and R. Thomas, Three-coloring triangle-free graphs on surfaces, in preparation.

[4] M. R. Garey and D. S. Johnson, Computers and intractability. A guide to the theory of NP-completeness, W. H. Freeman, San Francisco, 1979.

[5] J. Gimbel and C. Thomassen, Coloring graphs with fixed genus and girth, *Trans. Amer. Math. Soc.* **349** (1997), 4555–4564.

[6] H. Grötzsch, Ein Dreifarbensatz für dreikreisfreie Netze auf der Kugel, *Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg Math.-Natur. Reihe* **8** (1959), 109–120.

[7] Ł. Kowalik, Fast 3-coloring triangle-free planar graphs, Algorithms—ESA 2004, 436–447, *Lecture Notes in Comput. Sci.* **3221**, Springer, Berlin, 2004.

[8] Ł. Kowalik and M. Kurowski, Oracles for bounded length shortest paths in planar graphs, *ACM Trans. Algorithms* **2** (2006), 335–363.

[9] T. Nishizeki and N. Chiba, Planar graphs: theory and algorithms, *Ann. Discr. Math.* **32**, North-Holland, Amsterdam, 1988.

[10] N. Robertson, D. P. Sanders, P. D. Seymour and R. Thomas, The four-colour theorem, *J. Combin. Theory Ser. B* **70** (1997), 2-44.

[11] N. Robertson and P. D. Seymour, Graph Minors VII. Disjoint paths on a surface, *J. Combin. Theory Ser. B* **45** (1988), 212–254.

[12] W.-K. Shih and W.-L. Hsu, A new planarity test, *Theoret. Comp. Sci.* **223** (1999), 179–191.

[13] R. Thomas and B. Walls, Three-coloring Klein bottle graphs of girth five, *J. Combin. Theory Ser. B* **92** (2004), 115–135.

[14] C. Thomassen, Grötzsch's 3-color theorem and its counterparts for the torus and the projective plane, *J. Combin. Theory Ser. B* **62** (1994), 268–279.

[15] C. Thomassen, A short list color proof of Grotzsch's theorem, *J. Combin. Theory Ser. B* **88** (2003), 189–192.

[16] C. Thomassen, The chromatic number of a graph of girth 5 on a fixed surface, *J. Combin. Theory Ser. B* **87** (2003), 38-71.

[17] S. G. Williamson, Depth-first search and Kuratowski subgraphs, *J. Assoc. Comput. Mach.* **31** (1984), 681–693.