Dual Circumference and Collinear Sets

Vida Dujmović

School of Computer Science and Electrical Engineering, University of Ottawa, Canada
https://cglab.ca/~vida/
vida.dujmovic@uottawa.ca

Pat Morin

School of Computer Science, Carleton University, Canada
https://cglab.ca/~morin/
morin@scs.carleton.ca

Abstract

We show that, if an n-vertex triangulation T of maximum degree Δ has a dual that contains a cycle of length ℓ, then T has a non-crossing straight-line drawing in which some set, called a *collinear set*, of $\Omega(\ell/\Delta^4)$ vertices lie on a line. Using the current lower bounds on the length of longest cycles in 3-regular 3-connected graphs, this implies that every n-vertex planar graph of maximum degree Δ has a collinear set of size $\Omega(n^{0.8}/\Delta^4)$. Very recently, Dujmović et al. (SODA 2019) showed that, if S is a collinear set in a triangulation T then, for any point set $X \subset \mathbb{R}^2$ with $|X| = |S|$, T has a non-crossing straight-line drawing in which the vertices of S are drawn on the points in X. Because of this, collinear sets have numerous applications in graph drawing and related areas.

2012 ACM Subject Classification Mathematics of computing → Graph theory; Mathematics of computing → Extremal graph theory; Human-centered computing → Graph drawings

Keywords and phrases Planar graphs, collinear sets, untangling, column planarity, universal point subsets, partial simultaneous geometric drawings

Digital Object Identifier 10.4230/LIPIcs.SoCG.2019.29

Related Version A full version of this paper is available at https://arxiv.org/abs/1811.03427.

Funding This work was partly funded by NSERC and MRI.

Acknowledgements Much of this research took place during the Sixth Workshop on Order and Geometry held in Ciążeń, Poland, September 19–22, 2018. The authors are grateful to the organizers, Stefan Felsner and Piotr Micek, and to the other participants for providing a stimulating research environment.

1 Introduction

Throughout this paper, all graphs are simple and finite and have at least 4 vertices. For a planar graph G, we say that a set $S \subseteq V(G)$ is a *collinear set* if G has a non-crossing straight-line drawing in which the vertices of S are all collinear. A *plane graph* is a planar graph G along with a particular non-crossing drawing of G. The *dual* $G^\star$ of a plane graph G is the graph whose vertex set $V(G^\star)$ is the set of faces in G and in which $fg \in E(G^\star)$ if and only if the faces f and g of G have at least one edge in common. The *circumference*, $c(G)$, of a graph G is the length of the longest cycle in G. In Section 2, we prove the following:

▶ **Theorem 1.** *Let T be a triangulation of maximum degree Δ whose dual $T^\star$ has circumference ℓ. Then T has a collinear set of size $\Omega(\ell/\Delta^4)$.*

The dual of a triangulation is a 3-connected cubic planar graph. The study of the circumference of 3-connected cubic planar graphs has a long and rich history going back to at least 1884 when Tait [27] conjectured that every such graph is Hamiltonian. In 1946, Tait's conjecture was disproved by Tutte who gave a non-Hamiltonian 46-vertex example [28].

Repeatedly replacing vertices of Tutte's graph with copies of itself gives a family of graphs, $\langle G_i : i \in \mathbb{Z}\rangle$ in which G_i has $46 \cdot 45^i$ vertices and circumference at most $45 \cdot 44^i$. Stated another way, n-vertex members of the family have circumference $O(n^\alpha)$, for $\alpha = \log_{44}(45) < 0.9941$. The current best upper bound of this type is due to Grünbaum and Walther [18] who construct a 24-vertex non-Hamiltonian cubic 3-connected planar graph, resulting in a family of graphs in which n-vertex members have circumference $O(n^\alpha)$ for $\alpha = \log_{23}(22) < 0.9859$.

A series of results has steadily improved the lower bounds on the circumference of n-vertex (not necessarily planar) 3-connected cubic graphs. Barnette [5] showed that, for every n-vertex 3-connected cubic graph G, $c(G) = \Omega(\log n)$. Bondy and Simonovits [8] improved this bound to $e^{\Omega(\sqrt{\log n})}$ and conjectured that it can be improved to $\Omega(n^\alpha)$ for some $\alpha > 0$. Jackson [19] confirmed this conjecture with $\alpha = \log_2(1+\sqrt{5}) - 1 > 0.6942$. Billinksi et al. [6] improved this to the solution of $4^{1/\alpha} - 3^{1/\alpha} = 2$, which implies $\alpha > 0.7532$. The current record is held by Liu, Yu, and Zhang [22] who show that $\alpha > 0.8$.

It is known that any planar graph of maximum degree Δ can be triangulated so that the resulting triangulation has maximum degree $\lceil 3\Delta/2\rceil + 11$ [21]. This fact, together with Theorem 1 and the result of Liu, Yu, and Zhang [22], implies the following corollary:

► **Corollary 2.** *Every n-vertex triangulation of maximum degree Δ contains a collinear set of size $\Omega(n^{0.8}/\Delta^4)$.*

It is known that every planar graph G has a collinear set of size $\Omega(\sqrt{n})$ [9, 13] . Corollary 2 therefore improves on this bound for bounded-degree planar graphs and, indeed for the family of n-vertex planar graphs of maximum degree $\Delta \in O(n^\delta)$, with $\delta < 0.075$. For example, the triangulations dual to Grünbaum and Walther's construction have maximum degree $\Delta \in O(\log n)$. As discussed below, this implies that there exists n-vertex triangulations of maximum degree $O(\log n)$ whose largest collinear set has size $O(n^{0.9859})$. Corollary 2 implies that every n-vertex planar graph of maximum degree $O(\log n)$ has a collinear set of size $\Omega(n^{0.8})$.

Recently, Dujmović et al. [14] have shown that every collinear set is *free*. That is, for any planar graph G, any collinear set $S \subseteq V(G)$, and any set $X \subset \mathbb{R}^2$ with $|X| = |S|$, there exists a non-crossing straight-line drawing of G in which the vertices of S are drawn on the points of X. Because of this, collinear sets have immediate applications in graph drawing and related areas. For applications of Corollary 2, including untangling [11, 23, 29, 17, 20, 9, 12, 13, 25], column planarity [3, 15, 12, 13], universal point subsets [16, 1, 12, 13], and partial simultaneous geometric drawings [15, 4, 2, 7, 13] the reader is referred to Dujmović [13] and Dujmović et al. [14, Section 1.1]. Corollary 2 gives improved bounds for all of these problems for planar graphs of maximum $\Delta \in o(n^{0.075})$.

For example, it is known that every n-vertex planar geometric graph can be untangled while keeping some set of $\Omega(n^{0.25})$ vertices fixed [9] and that there are n-vertex planar geometric graphs that cannot be untangled while keeping any set of $\Omega(n^{0.4948})$ vertices fixed [10]. Although asymptotically tight bounds are known for paths [11], trees [17], outerplanar graphs [17], planar graphs of treewidth two [25], and planar graphs of treewidth three [12], progress on the general case has been stuck for 10 years due to the fact that the exponent 0.25 comes from two applications of Dilworth's Theorem. Thus, some substantially new idea appears to be needed. By relating collinear/free sets to dual circumference, the current paper presents an effective new idea. Indeed, Corollary 2 implies that every bounded-degree n-vertex planar geometric graph can be untangled while keeping $\Omega(n^{0.4})$ vertices fixed. Note that, even for bounded-degree planar graphs, $\Omega(n^{0.25})$ was the best previously-known lower bound.

Our work opens two avenues for further progress:

1. Lower bounds on the circumference of 3-regular 3-connected graphs is an active area of research. Indeed, the $\Omega(n^{0.8})$ lower bound of Liu, Yu, and Zhang [22] is less than a year old. Any further progress on these lower bounds will translate immediately to an improved bound in Corollary 2 and all its applications.
2. It is possible that the dependence on Δ can be removed from Theorem 1 and Corollary 2, thus making these results applicable to all planar graphs, regardless of maximum degree.

2 Proof of Theorem 1

Let G be a plane graph. We treat the vertices of G as points, the edges of G as closed curves, and the faces of G as closed sets (so that a face contains all the edges on its boundary and an edge contains both its endpoints). Whenever we consider subgraphs of G we treat them as having the same embedding as G. Similarly, if we consider a graph G' that is homeomorphic[1] to G then we assume that the edges of G' – each of which represents a path in G whose internal vertices all have degree 2 – inherit their embedding from the paths they represent in G.

Finally, if we consider the dual $G^\star$ of G then we treat it as a plane graph in which each vertex f is represented as a point in the interior of the face f of G that it represents. The edges of $G^\star$ are embedded so that an edge fg is contained in the union of the two faces f and g of G, it intersects the interior of exactly one edge of G that is common to f and g, and this intersection consists of a single point.

A *proper good curve* C for a plane graph G is a Jordan curve with the following properties:

1. *proper*: for any edge xy of G, C either contains xy, intersects xy in a single point (possibly an endpoint), or is disjoint from xy; and
2. *good*: C contains at least one point in the interior of some face of G.

Da Lozzo et al. [12] show that proper good curves define collinear sets:

▶ **Theorem 3.** *In a plane graph G, a set $S \subseteq V(G)$ is a collinear set if and only if there is a proper good curve for G that contains S.*

For a triangulation T, let $v(T)$ denote the size of a largest collinear set in T. We will show that, for any triangulation T of maximum degree Δ whose dual is $T^\star$, $v(T) = \Theta(c(T^\star)/\Delta^4)$ by relating proper good curves in T to cycles in $T^\star$.

As shown by Ravsky and Verbitsky [25, 24], the inequality $v(T) \le c(T^\star)$ is easy: If T is a triangulation that has a proper good curve C containing k vertices, then a slight deformation of C produces a proper good curve that contains no vertices. This curve intersects a cyclic sequence of faces $f_0, \dots, f_{k'-1}$ of T with $k' \ge k$. In this sequence, f_i and $f_{(i+1) \bmod k'}$ share an edge, for every $i \in \{0, \dots, k'-1\}$, so this sequence is a closed walk in the dual $T^\star$ of T. The properness of the original curve and the fact that each face of T is a triangle ensures that $f_i \ne f_j$ for any $i \ne j$, so this sequence is a cycle in $T^\star$ of length $k' \ge k$. Therefore, $c(T^\star) \ge v(T)$. From the result of Grünbaum and Walther described above, this implies that there are n-vertex triangulations T such that $v(T) = O(n^{0.9859})$.

The other direction, lower-bounding $v(T)$ in terms $c(T^\star)$ is more difficult. Not every cycle C of length ℓ in $T^\star$ can be easily transformed into a proper good curve containing a similar number of vertices in C. In the next section, we describe three parameters τ, ρ, and κ of a cycle C in $T^\star$ and show that C can always be transformed into a proper good curve containing $\Omega(\kappa)$ vertices of T.

[1] We say that a graph G' is homeomorphic to G if G' can be obtained from G by repeatedly contracting an edge of G that is incident to a degree-2 vertex.

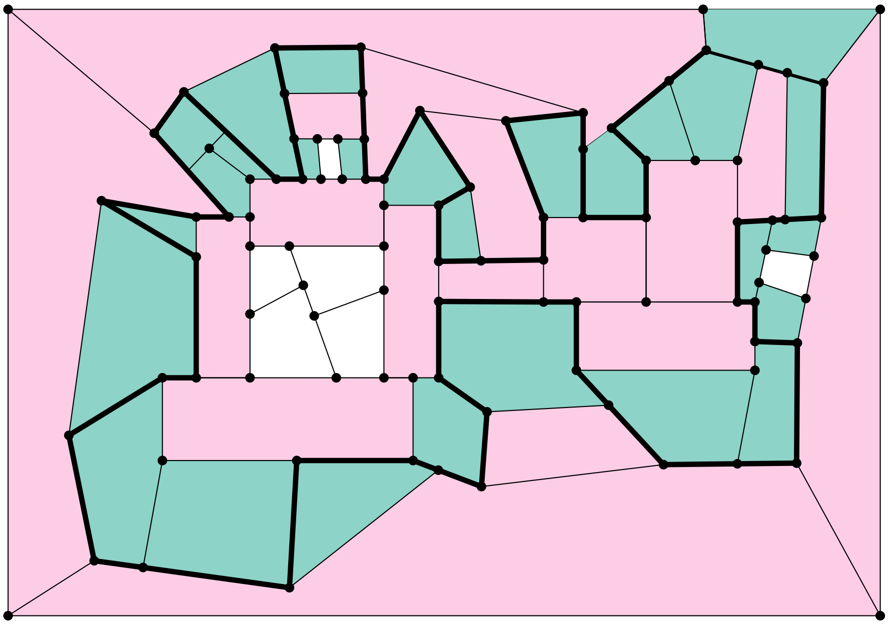

Figure 1 Faces of $T^\star$ that are pinched and caressed by C. C is bold, caressed faces are teal, pinched faces are pink, and untouched faces are unshaded.

2.1 Faces that are Touched, Pinched, and Caressed

Throughout the remainder of this section, T is a triangulation whose dual is $T^\star$ and C is a cycle in $T^\star$. Refer to Figure 1 for the following definitions. We say that a face f of $T^\star$

1. is *touched* by C if $f \cap C \neq \emptyset$;
2. is *pinched* by C if $f \cap C$ is a cycle or has more than one connected component; and
3. is *caressed* by C if it is touched but not pinched by C.

Since C is almost always the cycle of interest, we will usually say that a face f of $T^\star$ is touched, pinched, or caressed, without specifically mentioning C. We will frequently use the values τ, ρ, and κ to denote the number of faces of $T^\star$ in some region that are τouched, ρinched or κaressed. Observe that, since every face that is touched is either pinched or caressed, we have the identity $\tau = \rho + \kappa$.

▶ **Lemma 4.** *If C caresses κ faces of $T^\star$ then T has a proper good curve that contains at least $\kappa/4$ vertices so, by Theorem 3, $v(T) \geq \kappa/4$.*

Proof. Let F be the set of faces in $T^\star$ that are caressed by C. Each element $u \in F$ corresponds to a vertex of T so we will treat F as a set of vertices in T. Consider the subgraph $T[F]$ of T induced by F. The graph $T[F]$ is planar and has κ vertices. Therefore, by the 4-Colour Theorem [26], $T[F]$ contains an independent set $F' \subseteq F$ of size at least $\kappa/4$.

We claim that there is a proper good curve for T that contains all the vertices in F'. To see this, first observe that the cycle C in $T^\star$ already defines a proper good curve (that does not contain any vertices of T) that we also call C. We perform local modifications on C so that it contains all the vertices in F'.

For any vertex $u \in F'$, let $w_0, \ldots, w_{d-1}$ denote the neighbours of u in cyclic order. The curve C intersects some contiguous subsequence $uw_i, \ldots, uw_j$ of the edges adjacent to u. Since u is caressed, this sequence does not contain all edges incident to u. Therefore, the curve C crosses the edge $w_{i-1}w_i$, then crosses $uw_i, \ldots, uw_j$, and then crosses the edge w_jw_{j+1}. We

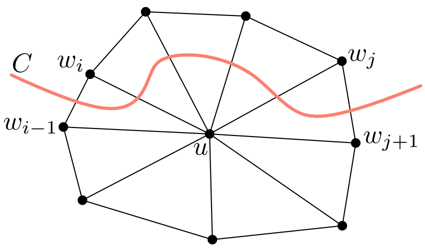

Figure 2 Transforming the dual cycle C into a proper good curve C' containing u.

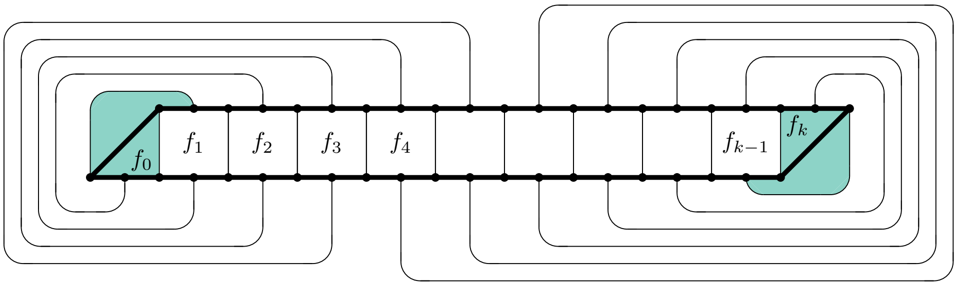

Figure 3 A Hamiltonian cycle C in $T^\star$ that caresses only four faces.

modify C by removing the portion between the first and last of these crossings and replacing it with a curve that contains u and is contained in the two triangles $w_{i-1}uw_i$ and w_juw_{j+1}. (See Figure 2.)

After performing this local modification for each $u \in F'$ we have a curve C' that contains every vertex $u \in F'$. All that remains is verify that C' is good and proper for T. That C' is good for T is obvious. That C' is proper for T follows from the following two observations: (i) C' does not contain any two adjacent vertices (since F' is an independent set); and (ii) if C' contains a vertex u, then it does not intersect the interior of any edge incident to u. ◀

Lemma 4 reduces our problem to finding a cycle in $T^\star$ that caresses many faces. It is tempting to hope that any sufficiently long cycle in $T^\star$ caresses many faces, but this is not true; Figure 3 shows that even a Hamiltonian cycle C in $T^\star$ may caress only four faces, two inside C and two outside of C. In this example, there is an obvious sequence of faces $f_0, \dots, f_k$, all contained in the interior of C where f_i shares an edge with f_{i+1} for each $i \in \{0, \dots, k-1\}$. The only faces caressed by C are the endpoints f_0 and f_k of this sequence.

Our strategy is to define a tree structure, T_0 on groups of faces contained in the interior of C and a similar structure, T_1 on groups of faces in the exterior of C. We will then show that every leaf of T_0 or T_1 contains a face caressed by C. In Figure 3, the tree T_0 is the path $f_0, \dots, f_k$ and, indeed, the leaves f_0 and f_k of this tree are caressed by C. After a non-trivial amount of analysis of the trees T_0 and T_1, we will eventually show that, if C does not caress many faces, then T_0 and T_1 have many nodes, but few leaves. Therefore T_0 and T_1 have many degree-2 nodes. This abundance of degree-2 nodes makes it possible to perform a surgery on C that increases the number of caressed faces. Performing this surgery repeatedly will then produce a curve C that caresses many faces.

A path $P = v_1, \dots, v_r$ in $T^\star$ is a *chord path* (for C) if $v_1, v_r \in V(C)$ and $v_2, \dots, v_{r-1} \notin V(C)$. Note that this definition implies that the interior vertices $v_2, \dots, v_{r-1}$ of P are either all contained in the interior of C or all contained in the exterior of C.

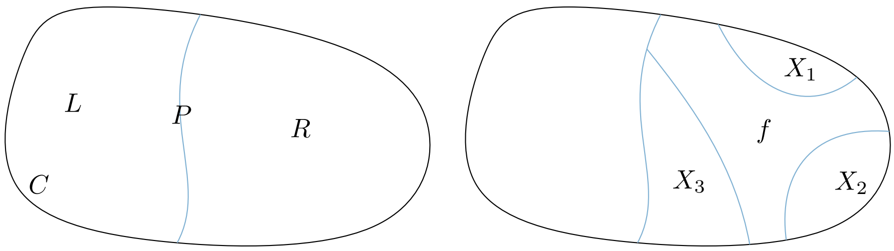

Figure 4 The proof of Lemma 5.

▶ **Lemma 5.** *Let P be a chord path for C and let L and R be the two faces of $P \cup C$ that each contain P in their boundary. Then R contains at least one face of $T^\star$ that is caressed by C.*

Proof. The proof is by induction on the number, t, of faces of $T^\star$ contained in R. If $t = 1$, then R is a face of $T^\star$ and it is caressed by C.

If $t > 1$, then consider the face f of $T^\star$ that is contained in R and has the first edge of P on its boundary. Refer to Figure 4. Since $t > 1$, $X = R \setminus f$ is non-empty. The set X may have several connected components $X_1, \ldots, X_k$, but each X_i has a boundary that contains a chord path P_i for C. We can therefore apply induction on P_1 (or any P_i) using $R = X_1$ in the inductive hypothesis. ◀

2.2 Auxilliary Graphs and Trees: H, $\tilde{H}$, T_0, and T_1

Refer to Figure 5. Consider the auxilliary graph H with vertex set $V(H) \subseteq V(T^\star)$ and whose edge set consist of the edges of C plus those edges of $T^\star$ that belong to any face pinched by C. Let $v_0, \ldots, v_{r-1}$ be the clockwise cyclic sequence of vertices on some face f of $T^\star$ that is pinched by C. We identify three kinds of vertices that are *special* with respect to f:

1. A vertex v_i is special of *Type A* if $v_{i-1}v_i$ is an edge of C and v_iv_{i+1} is not an edge of C.
2. A vertex v_i is special of *Type B* if $v_{i-1}v_i$ is not an edge of C and v_iv_{i+1} is an edge of C.
3. A vertex v_i is special of *Type Y* if v_i not incident to any edge of C and v_i has degree 3 in H.

We say that a chord path $v_i, \ldots, v_j$ is a *keeper* with respect to f if v_i is special of Type A, v_j is special of Type B, and none of $v_{i+1}, \ldots, v_{j-1}$ are special. We let $\tilde{H}$ denote the subgraph of H containing all the edges of C and all the edges of all paths that are keepers with respect to some pinched face f of $T^\star$.

It is worth emphasizing at this point that, by definition, every keeper is entirely contained in the boundary of at least one face f of $T^\star$. This property will be useful shortly.

Let $\tilde{H}'$ denote the graph that is homeormophic to $\tilde{H}$ but does not contain any degree 2 vertices. That is, $\tilde{H}'$ is the minor of $\tilde{H}$ obtained by repeatedly contracting an edge incident a degree-2 vertex. The graph $\tilde{H}'$ naturally inherits an embedding from the embedding of $\tilde{H}$. This embedding partitions the edges of $\tilde{H}'$ into three sets:

1. The set B of edges that are contained in (the embedding of) C;
2. The set E_0 of edges whose interiors are contained in the interior of (the embedding of) C; and
3. The set E_1 of edges whose interiors are contained in the exterior of (the embedding of) C.

Observe that, for each $i \in \{0, 1\}$, the graph H_i whose edges are exactly those in $B \cup E_i$ is outerplanar, since all vertices of H_i are on a single face, whose boundary is C. Let ${H_i}^\star$ be dual of H_i and let T_i be the subgraph of ${H_i}^\star$ whose edges are all those dual to the edges of E_i. From the outerplanarity of H_i, it follows that T_i is a tree.

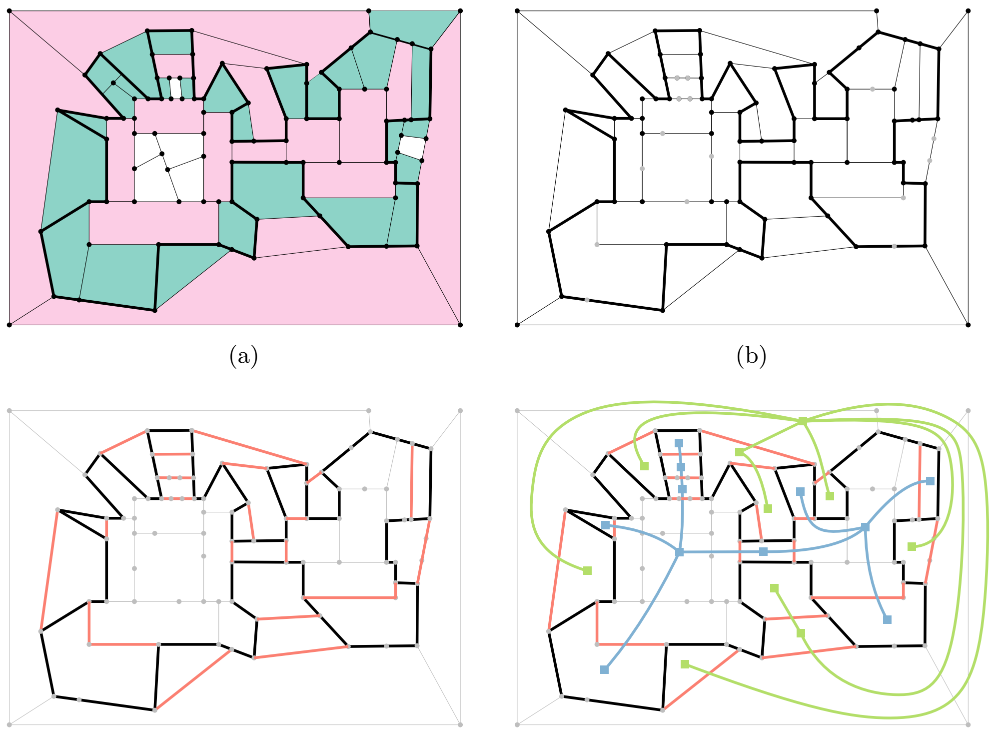

(a) (b) (c) (d)

Figure 5 (a) the cycle C in $T^\star$ with faces classified as pinched or caressed; (b) the auxilliary graph H; (c) the auxilliary graph $\tilde{H}$ with keeper paths highlighted; (d) the trees T_0 and T_1.

Each vertex of T_i corresponds to a face of $\tilde{H}$. From this point onwards, we will refer to the vertices of T_i as *nodes* to highlight this fact, so that a node u of T_i is synonymous with the subset of $\mathbb{R}^2$ contained in the corresponding face of $\tilde{H}$. In the following, when we say that a node u of T_i contains a face f of $T^\star$ we mean that f is one of the faces of $T^\star$ whose union makes up u. The degree, δ_u of any node u in T_i is exactly equal to the number of keeper paths on the boundary of u.

The following lemma allows us to direct our effort towards proving that one of T_0 or T_1 has many leaves.

▶ **Lemma 6.** *Each leaf u of T_i contains at least one face of $T^\star$ that is caressed by C.*

Proof. The edge of T_i incident to u corresponds to a chord path P. The graph $P \cup C$ has two faces with P on its boundary, one of which is u. The lemma now follows immediately from Lemma 5, with $R = u$. ◀

We will make use of the following well-known property of 3-connected plane graphs.

▶ **Lemma 7.** *If T has $n \geq 4$ vertices then any two faces of $T^\star$ share at most one edge.*

▶ **Lemma 8.** *Let u be a node of T_i and let ρ_u, κ_u, and δ_u denote the number of pinched faces of $T^\star$ in u, the number of caressed faces of $T^\star$ in u, and the degree of u in T_i, respectively. Then $\rho_u \leq 2(\kappa_u + \delta_u)$.*

Before proving Lemma 8, we point out that the leading constant 2 is tight. Figure 6 shows an example in which all $\rho_u = 2k+1$ pinched faces of $T^\star$ are contained in a single (pink) node u of T_0 that contains $\kappa_u = 0$ caressed faces and has degree $\delta_u = k+2$.

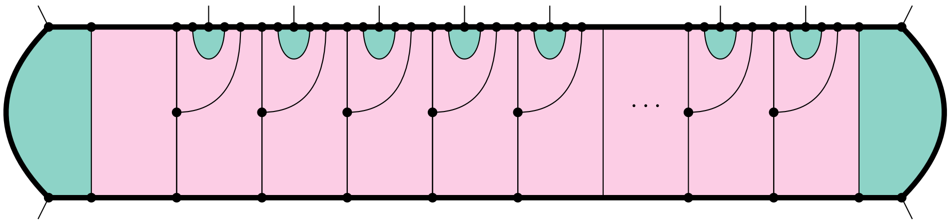

Figure 6 An example showing the tightness of Lemma 8.

Proof. The proof is a discharging argument. We assign each pinched face in u a single unit of charge, so that the total charge is ρ_u. We then describe a discharging procedure that preserves the total charge. After executing this procedure, pinched faces in u have no charge, each caressed face in u has charge at most 2, and each keeper path in u has charge at most 2. Since there is a bijection between keeper paths in u and edges of T_i incident to u, this proves the result.

We now describe the discharging procedure, which is recursive and takes as input a chord path P that partitions u into two parts L and R. We require as a precondition that there are $m \geq 1$ pinched faces of $T^\star$ in L, each of which contains at least one edge of P and such that every edge of P is contained in at least one of these faces. During a recursive call, P may have a charge $c \in \{0, 1, 2\}$. This charge will be at most 1 if $m > 1$, but can be 2 if $m = 1$.

To initialize the discharging procedure, we choose an arbitrary pinched face f contained in u. The face f begins with one unit of charge and has $r \geq 2$ chord paths $P_1, \ldots, P_r$ on its boundary. We move the charge from f onto P_1 and apply the recursive procedure to P_1, with a charge of 1 (with L being the component of $u \setminus P_1$ that contains f). We then recursively apply the discharging procedure on each of $P_2, \ldots, P_r$ with a charge of 0.

Next we describe each recursive step, during which we are given P with some charge $c \in \{0, 1, 2\}$. There are several cases to consider (see Figure 7):

1. R contains no face of $T^\star$ that is pinched by C. If R is empty, then P is a keeper path, in which case we leave a charge of c on it and we are done. Otherwise R is non-empty and Lemma 5 ensures that R contains at least one caressed face f. We move the charge from P onto f and we are done.
2. R contains a face f that is pinched by C and that shares at least one edge with P. We consider three subcases:
   a. f contains neither endpoint of P. In this case, $R \setminus f$ has at least three connected components, A, B, and $X_1, \ldots, X_k$, where A and B each contain an endpoint of P and each X_i has a chord path P_i in common with f. We recurse on each of these components so that each of these components takes the place of R in the recursion. When recursing on A we take one unit of charge from P (if needed) and place it on A's chord path. When recursing on B we take the second unit of charge from P (if needed) and place it on B's chord path. When recursing on X_1 we move the unit of charge from f to P_1. When recursing on $X_2, \ldots, X_k$ we use no charge on $P_2, \ldots, P_k$.
   b. f contains exactly one endpoint of P. In this case, $R \setminus f$ has one connected component A that contains an endpoint of P and one or more connected components $X_1, \ldots, X_k$ where each X_i has a chord path P_i on the boundary of f. The path P has a charge $c \leq 2$. When recursing on X_1 we assign all of P's charge to the chord path P_1, which is contained in the single pinched face f. When recursing on A we move the single unit of charge from f to the chord path of A.

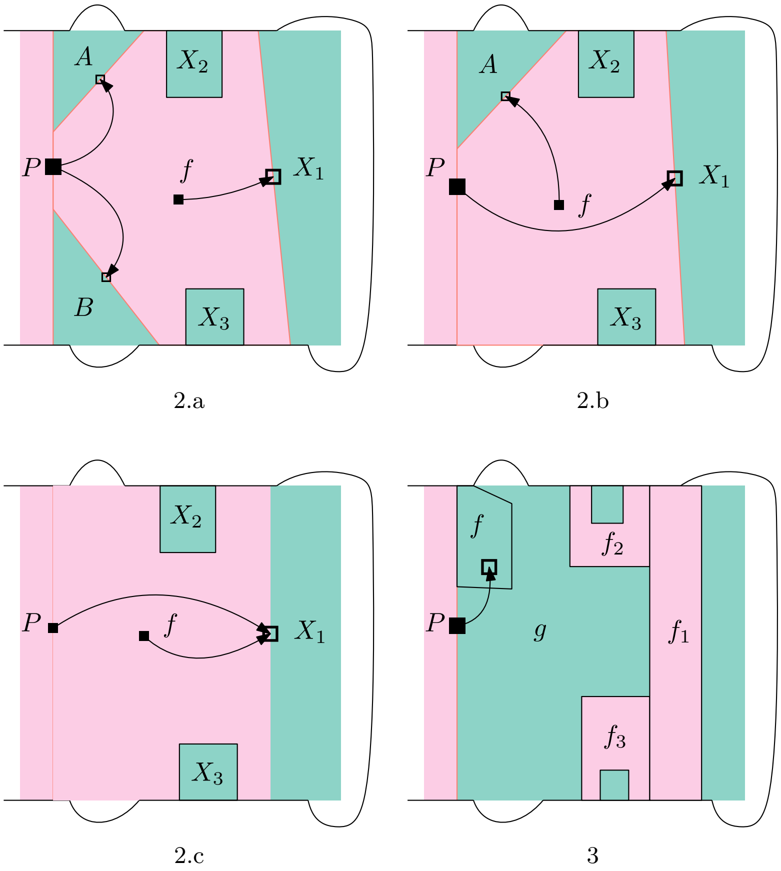

Figure 7 Discharging steps in the proof of Lemma 8.

c. f contains both endpoints of P. We claim that, in this case, P must be on the boundary of more than one pinched faces in L, otherwise P would be a keeper path. To see this, observe that the face f contains both the first edge e_1 and last edge e_2 of P. If $e_1 = e_2$ because P is a single edge, then it is certainly a keeper, which is not possible. Otherwise, by Lemma 7, e_1 and e_2 are on the boundary of two different faces in L. By assumption, both of these faces are pinched by C.
Therefore P has at most one unit of charge assigned to it. Now, $R \setminus f$ has one or more connected components $X_1, \ldots, X_k$ sharing chord paths $P_1, \ldots, P_k$ with f on which we recurse. When recursing on X_1 we move the charge from P and the charge from f to P_1. When recursing on the remaining X_i, $i \in \{2, \ldots, k\}$ we assign no charge to P_i.

3. R contains at least one pinched face, but no pinched face in R shares an edge with P. In this case, consider the face g of H that is contained in R and has P on its boundary. By definition, g contains no pinched faces of $T^\star$, but g is touched by C, so g contains at least one caressed face[2] f of $T^\star$. We move the c units of charge from P onto f.
Now, R still contains one or more pinched faces $f_1, \ldots, f_k$, where each f_i shares part of a chord path P_i with g. On each such face f_i, we run the initialization procedure described above except that we recurse only on the chord paths of f_i that do not share edges with g. i.e., we do not recurse on the chord path P_i.

This completes the description of the discharging procedure, and the proof. ◀

[2] In fact g contains at least two caressed faces, one for each endpoint of P.

2.3 Bad Nodes

We say that a node of T_i is *bad* if it has degree 2 and contains no face of $T^\star$ that is caressed by C. We now move from studying individual nodes of T_0 and T_1 to studying global quantities associated with T_0 and T_1. From this point on, for each $i \in \{0, 1\}$,

1. τ_i, ρ_i, and κ_i refer the total numbers of faces contained in nodes of T_i that are touched, pinched, and caressed by C, respectively;
2. n_i refers to the number of nodes of T_i;
3. $\delta_i = 2(n_i - 1)$ is the total degree of all nodes in T_i; and
4. b_i is the number of bad nodes in T_i.

▶ **Lemma 9.** *If $\kappa_i \le \tau_i/6$ then $n_i \ge \tau_i/8$.*

Proof. From Lemma 8 we know $\rho_i \le 2(\kappa_i + \delta_i)$, so

$$\tau_i = \kappa_i + \rho_i \le 3\kappa_i + 2\delta_i = 3\kappa_i + 4(n_i - 1) \le \tau_i/2 + 4n_i \ ,$$

and reorganizing the left- and right-hand sides gives the desired result. ◀

▶ **Lemma 10.** *For any $0 < \epsilon < 1$, if $b_i \le (1 - \epsilon)n_i$, then $\kappa_i = \Omega(\epsilon\tau_i)$.*

Proof. Partition the nodes of T_i into the following sets:

1. the set B of bad nodes;
2. the set N_1 of leaves;
3. the set $N_{\ge 3}$ of nodes having degree at least 3;
4. the set N_2 of nodes having degree 2 that are not bad.

$$\begin{aligned}
b_i &= n_i - |N_1| - |N_{\ge 3}| - |N_2| \\
&> n_i - 2|N_1| - |N_2| && \text{since } |N_1| > |N_{\ge 3}| \\
&\ge n_i - 2\kappa_i - |N_2| && \text{(since, by Lemma 6, } \kappa_i \ge |N_1|) \\
&\ge n_i - 3\kappa_i && \text{(since each node in } N_2 \text{ contains a caressed face)}
\end{aligned}$$

Thus, we have

$$n_i - 3\kappa_i \le b_i \le (1 - \epsilon)n_i$$

and rewriting gives

$$\kappa_i \ge \epsilon n_i/3 \ . \tag{1}$$

If $\kappa_i \ge \tau_i/6$, then the proof is complete. On the other hand, if $\kappa_i \le \tau_i/6$ then, by Lemma 9, $n_i \ge \tau_i/8$. Combining this with (1) gives

$$\kappa_i \ge \epsilon n_i/3 \ge \epsilon\tau_i/24 = \Omega(\epsilon\tau_i) \ .$$ ◀

2.4 Interactions Between Bad Nodes

We have now reached a point in which we know that the vast majority of nodes in T_0 and T_1 are bad nodes, otherwise Lemma 10 implies that a constant fraction of the faces touched by C are caressed by C. At this point, we are ready to study interactions between bad nodes of T_0 and bad nodes of T_1. The proof of the following lemma is ommitted due to space constraints but can be found in the full version of the paper.

▶ **Lemma 11.** *If u is a bad node then there is a single face f of $T^\star$ that is contained in u and that contains all edges of $C \cap u$.*

The following lemma shows that a bad node u in T_0 and a bad node w in T_1 share at most one edge of C.

▶ **Lemma 12.** *If nodes u in T_0 and w in T_1 are bad nodes that share at least one edge of C, then u and w share exactly one edge of C.*

Proof. Suppose u and w share two edges e_1 and e_2 of C. Then, by Lemma 11, there is a common face f_u in u that contains e_1 and e_2, Similarly, there is a common face f_w contained in w that contains both e_1 and e_2. But this contradicts Lemma 7. ◀

▶ **Lemma 13.** *If u and w are bad nodes of T_i sharing a common chord path P, then P is a single edge.*

Proof. By Lemma 11, u and w have the first edge of P in common and the last edge of P in common. Lemma 7 therefore implies that the first and last edge of P are the same, so P has only one edge. ◀

2.5 Really Bad Nodes

At this point we will start making use of the assumption that the triangulation T has maximum degree Δ, which is equivalent to the assumption that each face of $T^\star$ has at most Δ edges on its boundary.

▶ **Observation 14.** *If T has maximum degree Δ and C has length ℓ, then the number of faces τ of $T^\star$ touched by C is at least $2\ell/\Delta$. At least ℓ/Δ of these faces are in the interior of C and at least ℓ/Δ of these faces are in the exterior of C.*

Proof. Orient the edges of C counterclockwise so that, for each edge e of C, the face of $T^\star$ to the left of e is in C's interior and the face of $T^\star$ to the right of e is in C's exterior. Each face of $T^\star$ has at most Δ edges. Therefore, the number of faces to the right of edges in C is at least ℓ/Δ. The same is true for the number of faces of $T^\star$ to the left of edges in C. ◀

For a node u of T_i, we define $N(u)$ as the set of nodes in T_0 and T_1 (excluding u) that share an edge of $T^\star$ with u. Note that $N(u)$ contains the neighbours of u in T_i as well as nodes of T_{1-i} with which u shares an edge of C. We say that a node u is *really bad* if u and all nodes in $N(u)$ are bad. The proof of the following lemma – which is similar to that of Lemma 10 – is ommitted due to space constraints.

▶ **Lemma 15.** *For every sufficiently small $0 < \alpha < 1/2$, if T has maximum degree Δ, C has length ℓ, and the number κ, of faces caressed by C is at most $\alpha\ell/\Delta^2$, then the number of really bad nodes in T_0 is at least $n_0 - O(\alpha n_0)$.*

For a node u of T_i, we define $N^0(u) = \{u\}$ and, for any $r \in \mathbb{N}$, we define $N^r(u) = \bigcup_{w \in N^{r-1}(u)} N(w)$. We say that a node u in T_i is *really^r bad* if u is bad and all nodes in $N^r(u)$ are bad. The proof of the following lemma is a straightforward generalization of the proof of Lemma 15.

▶ **Lemma 16.** *For any constant $i \in \mathbb{N}$ and every sufficiently small $0 < \alpha < 1/2$, if T has maximum degree Δ, C has length ℓ, and the number, κ, of faces caressed by C is at most $\alpha\ell/\Delta^{i+1}$, then the number of really^i bad nodes in T_0 is at least $n_0 - O(\alpha n_0)$.*

For our purposes, it will be sufficient to work with bad ($i = 0$), really bad ($i = 1$), and really really bad ($i = 2$) nodes.

2.6 Tree/Cycle Surgery

We summarize the situation so far. We are left with the case where C has length ℓ and therefore touches $\Omega(\ell/\Delta)$ faces. To complete the proof of Theorem 1 we must deal with the situation where C caresses $o(\ell/\Delta^4)$ faces and therefore each of T_0 and T_1 has $o(\ell/\Delta^4)$ leaves (Lemma 6), $\Omega(\ell/\Delta)$ nodes (Lemma 9), and the fraction of really really bad nodes in T_0 and T_1 is $1 - O(1/\Delta^2)$ (Lemma 16).

To handle cases like these, the only option is to perform surgery on the cycle C to increase the number of caressed nodes. In particular, our strategy is to perform modifications to C that increase the number of faces caressed by C. At this point we are ready to complete the proof of Theorem 1.

Proof of Theorem 1. By Lemma 4, it suffices to prove the existence of a cycle C in $T^\star$ that caresses $\Omega(\ell/\Delta^4)$ faces. We begin by applying Lemma 16 with $i = 2$ and $\alpha = \epsilon/\Delta$. For sufficiently small, but constant, ϵ, Lemma 16 implies that $\kappa = \Omega(\ell/\Delta^4)$ or the number of nodes in T_0 that are not really really bad is at most $O(\epsilon n_0/\Delta)$. In the former case, C caresses $\Omega(\ell/\Delta^4)$ faces of $T^\star$ and we are done.

In the latter case, consider the forest obtained by removing all nodes of T_0 that are not really really bad. This forest has $(1 - O(\epsilon/\Delta))n_0$ nodes. We claim that it also has $O(\epsilon n_0/\Delta)$ components. To see why this is so, let L be the set of leaves in T_i and let S be the set of non-leaf nodes in T_i that are not really really bad. Observe that it is sufficient to upper bound the number, k of components in $T_i - S$.

We have $|L| \le |S| + |L| = O(\epsilon n_0/\Delta)$ and $k = \sum_{u\in S}(\deg_{T_i}(u) - 1)$. Furthermore,

$$O(\epsilon n_0/\Delta) \ge |L| \ge \sum_{u\in S}(\deg_{T_i}(u) - 2) = |S| + \sum_{u\in S}(\deg_{T_i}(u) - 1) = |S| + k \ .$$

Therefore $k \le |S| + k = O(\epsilon n_0/\Delta)$, as claimed.

Thus the forest obtained by removing all really really bad nodes from T_i has at most $O(\epsilon n_0/\Delta)$ components, each of which is a path. At least one of these paths contains $\Omega(\Delta/\epsilon)$ nodes. In particular, for a sufficiently small constant ϵ, one of these components, X, has at least 5Δ nodes.

Consider some node u in X, and let C_a and C_b be the two components of $u \cap C$. Observe that $T_1[N(u)]$ consists of two paths $a_1, \dots, a_r$ and $b_1, \dots, b_s$ of bad nodes where each $a_1, \dots, a_r$ contains an edge of C_a and each of $b_1, \dots, b_r$ contains an edge of C_b. Note that it is possible that $a_i = b_j$ for some values of i and j, but everything stated thus far, and subsequently, is still true. It follows from Lemma 12 that among any sequence of Δ consecutive nodes in X, at least one node has $r \ge 2$ and therefore $|N(u)| \ge 5$. Let u be any such node that is not among the first 2Δ or last 2Δ nodes of X. Such a u always exists because X contains at least 5Δ nodes.

Let $x_0 = u$. We now describe some of the nodes in the vicinity of u (refer to Figure 8):

1. there is a path $x_{2\Delta}, \dots, x_1, x_0, y_1, \dots, y_{2\Delta}$ in T_0 consisting entirely of really really bad nodes.
2. some really bad node a_1 of T_1 shares an edge with each of $x_0, \dots, x_i$ for some $i \in \{1, \dots, \Delta - 4\}$.
3. some really bad node a_2 of T_1 shares an edge with a_1 and and edge with x_0.
4. some really bad node a_0 of T_1 shares an edge with a_1 and with each of $x_i, \dots, x_{i+j}$ for some $j \in \{0, \dots, \Delta - 4\}$.

The surgery we perform focuses on the nodes u and a_1. Consider the two components of $C \cap a_1$. One of these components, p, shares an edge with u. By Lemma 12, the other component, q, does not share an edge with u. Imagine removing u from T_0, thereby separating

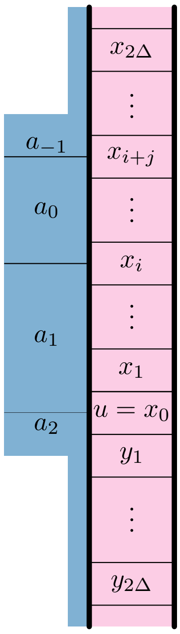

Figure 8 Nodes in the vicinity of $u = x_0$.

T_0 into a component T_x containing x_1 and a component T_y containing y_1. Equivalently, one can think of removing the edges of u from C separating C into two paths C_x and C_y on the boundary of T_x and T_y, respectively. We distinguish between two major cases (see Figure 9):

1. $q \subset C_x$. In this case, we punt to Case 2. By Lemma 13 $a_1 - C$ consists of two edges and exactly one of these edges, e, is not incident to u. Instead, e is incident to x_i. We set $u' = x_i$, $x'_1 = x_{i-1}$, $y'_1 = x_{i+1}$, and $a'_1 = a_1$. Observe that a'_1 connects the two components of $T_0 - u'$ and shares edges with u' and x'_1. This is exactly the situation considered in Case 2.
2. $q \subset C_y$. At this point it is helpful to think of T_0, T_1, and C as a partition of $\mathbb{R}^2$, where nodes of T_0 are coloured red, nodes of T_1 are coloured blue and C is the (purple) boundary between red and blue. To describe our modifications of C, we imagine changing the colours of nodes. The effect that such a recolouring has on C is immediately obvious: It produces a 1-dimensional set C' that contains every (purple) edge contained in the red-blue boundary. The set C' is a collection of vertices and edges of $T^\star$. Therefore, if C' is a simple cycle, then C' defines a new pair of trees T'_0 and T'_1.
   Refer to the right two thirds of Figure 9 for a simple (and misleading) example of what follows. For a full example, refer to Figure 10. The surgery we perform, recolours $x_0, x_1, \ldots, x_{i-1}$ blue and recolours a_1 red. Observe that, because $q \subset C_y$ and p contain an edge of x_1, this implies that the red subset of $\mathbb{R}^2$ is simply-connected and its boundary C' is a simple cycle consisting of edges of T^*. The new trees T'_0 and T'_1 are therefore well defined. We now make two claims that will complete our proof.

   ▷ Claim 17. For each $i \in \{0, 1\}$, and every node of w T_i that is not bad, $C \cap w = C' \cap w$. (Equivalently, for every face f of $T^\star$ that is not a bad node of T_0 or T_1, $C \cap f = C' \cap f$.)

   ▷ Claim 18. The face a_0 is caressed by C'.

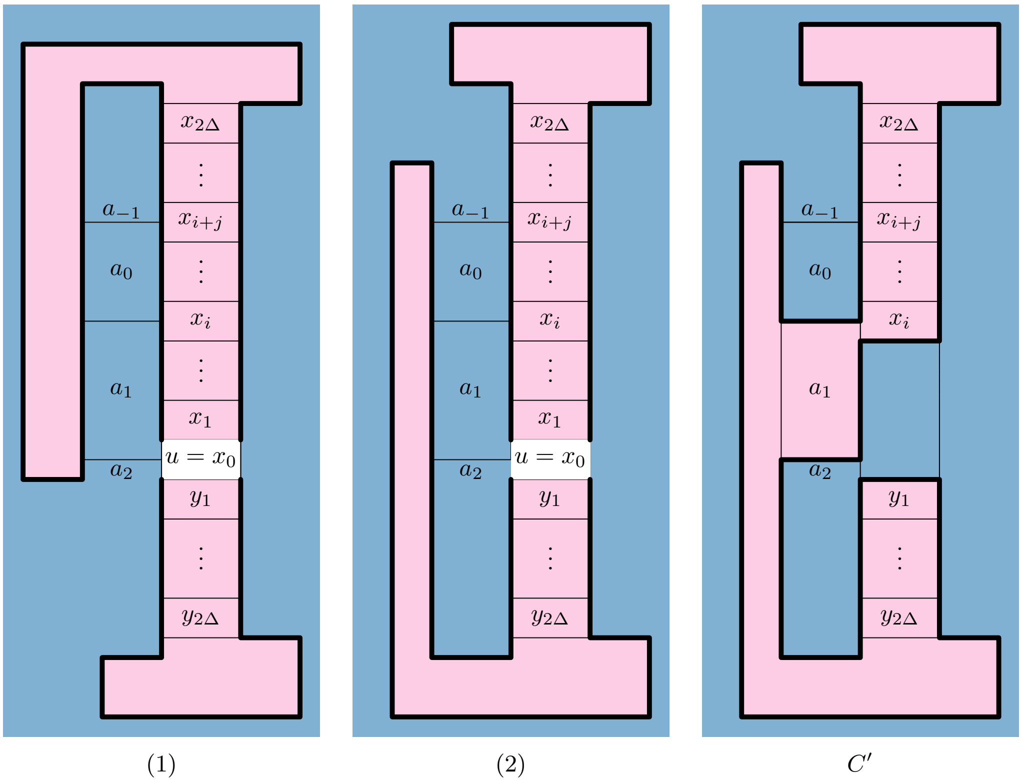

Figure 9 Cases 1 and 2 in the proof of Theorem 1 and the surgery performed in Case 2.

These two claims complete the proof because, together, they imply that C' caresses one more node than C. Indeed, by definition, C did not caress any faces belonging to bad nodes. Therefore, the first claim implies that the faces of $T^\star$ caressed by C' are a superset of those caressed by C. The face a_0 is a bad node of T_i so it is not caressed by C but the second claim states that it is caressed by C'. Therefore C' caresses at least one more face than C.

This surgery recolours at most $\Delta - 2 \leq \Delta$ nodes of T_0 and T_1, so the difference in length between C and C' is at most Δ^2. If we start with a cycle C of length ℓ, then we can perform this surgery at least $\ell/(4\Delta^2)$ times before the length of C decreases to less than $\ell' = \ell/2$. If at some point during this process, we are no longer able to perform this operation, it is because C caresses $\Omega(\ell'/\Delta^4) = \Omega(\ell/\Delta^4)$ faces of $T^\star$ and we are done. By the end of this process, the number of faces caressed by C is at least $\ell/(4\Delta^2) \in \Omega(\ell/\Delta^2) \subset \Omega(\ell/\Delta^4)$ and we are also done.

Thus, all that remains is to prove Claim 17 and Claim 18.

To prove Claim 17 we observe that C and C' differ only on the boundaries of nodes that are recoloured. Thus, it is sufficient to show that all nodes in $R = \cup\{N(v) : v \in \{x_0, \ldots, x_{i-1}, a_1\}$ are bad. But this is immediate since $x_0, \ldots, x_{i-1}$ are really really bad and $a_1 \in N(x_0)$, so a_1 is bad. Since every node in R share an edge with at least one of $\{x_0, \ldots, x_{i-1}, a_1\}$, every node in R is therefore bad, as required.

To prove Claim 18 we consider the boundary of the face a_0 of T^* after the recolouring operation. This boundary consists of, in cyclic order:

a. An edge shared between a_0 and a_1. This edge is in C' since a_0 is in T'_1 and a_1 is in T'_0.

b. A path of edges shared with $x_i, \ldots, x_{i+j}$. The nodes $x_i, \ldots, x_{i+j}$ are in T_0 and are distinct from $x_0, \ldots, x_{i-1}$, so these nodes are in T'_0. Therefore, this part of the bounary of a_0 is contained in C'.

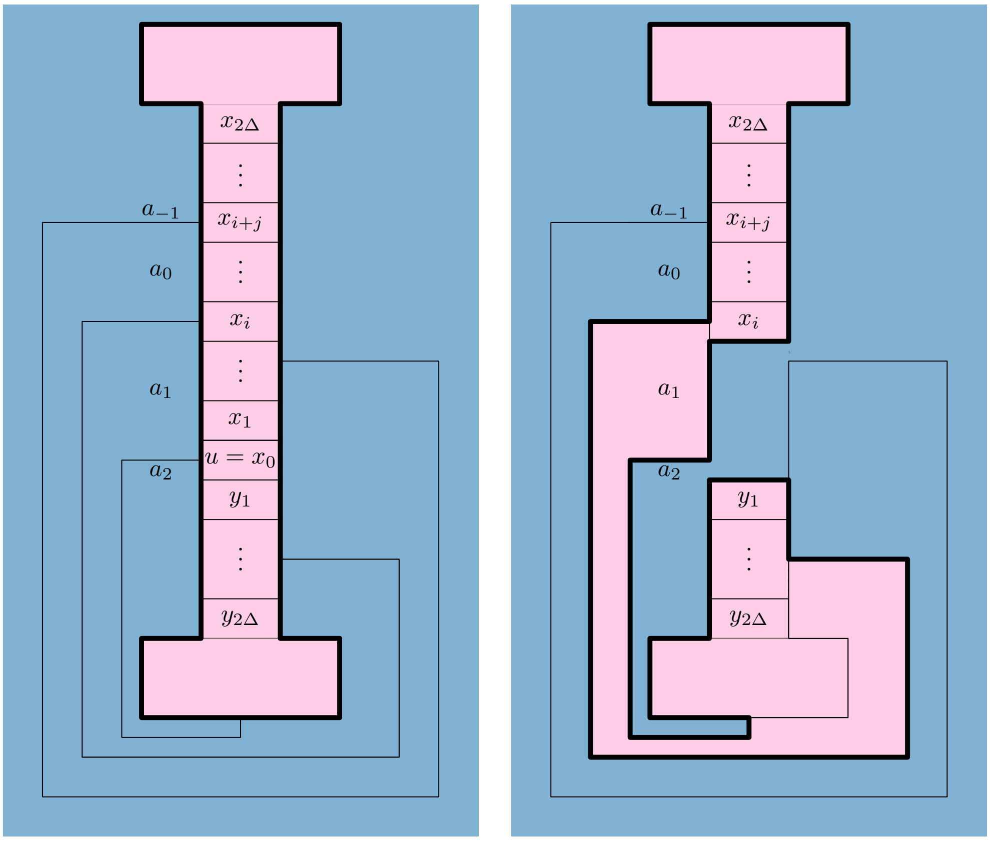

Figure 10 Performing surgery on C to obtain C' that caresses a_0.

c. An edge shared between a_0 and another node $a_{-1} =\neq a_1$ of T_1. The faces of a_{-1} are in T_1' because a_1 is the only face that moves from T_1 to T_0'. (T_1' is the only face whose colour goes from blue to red.)

d. A path of edges that contains at least one edge of C_y. If we fix an embedding in which the outer face is some face of T_1 other than a_0, then this path contains a portion of C that is traversed in clockwise order. By Lemma 12 This path does not contain any edge of x_i. Furthermore, this path does not contain any edges of $x_0, x_1, \ldots, x_{i-1}$ that are not on the outer face of $T_0 \cup a_1$. Therefore, this path consists of a (possibly empty) sequence of edges that are shared with $x_0, \ldots, x_{i-1}$ followed by a sequence of edges from C_y. The former part of this path is shared with nodes in T_1', so these edges are not in C'. The latter part of this path is shared with nodes in T_y, which are all contained contained in T_0'.

Therefore the intersection $C' \cap a_0$ consists of one connected component so a_0 is caressed by C'. ◀

3 Discussion

It remains an open problem to eliminate the dependence of our results on the maximum degree, Δ, of T. The next significant step is to resolve the following conjecture:

▶ **Conjecture 19.** *If T is a triangulation whose dual $T^\star$ has a cycle of length ℓ, then $T^\star$ has a cycle that caresses $\Omega(\ell)$ faces. (Therefore, by Lemma 4 and Theorem 3, T has a collinear set of size $\Omega(\ell)$.)*

References

1 Patrizio Angelini, Carla Binucci, William S. Evans, Ferran Hurtado, Giuseppe Liotta, Tamara Mchedlidze, Henk Meijer, and Yoshio Okamoto. Universal Point Subsets for Planar Graphs. In Kun-Mao Chao, Tsan-sheng Hsu, and Der-Tsai Lee, editors, *Algorithms and Computation - 23rd International Symposium, ISAAC 2012, Taipei, Taiwan, December 19-21, 2012. Proceedings*, volume 7676 of *Lecture Notes in Computer Science*, pages 423–432. Springer, 2012. doi:10.1007/978-3-642-35261-4_45.

2 Patrizio Angelini, William S. Evans, Fabrizio Frati, and Joachim Gudmundsson. SEFE without mapping via large induced outerplane graphs in plane graphs. *Journal of Graph Theory*, 82(1):45–64, 2016. doi:10.1002/jgt.21884.

3 Luis Barba, William Evans, Michael Hoffmann, Vincent Kusters, Maria Saumell, and Bettina Speckmann. Column planarity and partially-simultaneous geometric embedding. *J. Graph Algorithms Appl.*, 21(6):983–1002, 2017. doi:10.7155/jgaa.00446.

4 Luis Barba, Michael Hoffmann, and Vincent Kusters. Column planarity and partial simultaneous geometric embedding for outerplanar graphs. In *Abstracts of the 31st European Workshop on Computational Geometry (EuroCG)*, pages 53–56, 2015.

5 David Barnette. Trees in polyhedral graphs. *Canadian Journal of Mathematics*, 18:731–736, 1966.

6 Mark Bilinski, Bill Jackson, Jie Ma, and Xingxing Yu. Circumference of 3-connected claw-free graphs and large Eulerian subgraphs of 3-edge connected graphs. *J. Combin. Theory Ser. B*, 101:214–236, 2011.

7 T. Bläsius, S. G. Kobourov, and I. Rutter. Simultaneous embedding of planar graphs. In Roberto Tamassia, editor, *Handbook of Graph Drawing and Visualization*, pages 349–381. Chapman and Hall/CRC Press, 2013.

8 J.A. Bondy and M. Simonovits. Longest cycles in 3-connected 3-regular graphs. *Canadian Journal of Mathematics*, 32:987–992, 1980.

9 Prosenjit Bose, Vida Dujmović, Ferran Hurtado, Stefan Langerman, Pat Morin, and David R. Wood. A Polynomial Bound for Untangling Geometric Planar Graphs. *Discrete & Computational Geometry*, 42(4):570–585, 2009. doi:10.1007/s00454-008-9125-3.

10 Javier Cano, Csaba D. Tóth, and Jorge Urrutia. Upper Bound Constructions for Untangling Planar Geometric Graphs. *SIAM J. Discrete Math.*, 28(4):1935–1943, 2014. doi:10.1137/130924172.

11 Josef Cibulka. Untangling Polygons and Graphs. *Discrete & Computational Geometry*, 43(2):402–411, 2010. doi:10.1007/s00454-009-9150-x.

12 Giordano Da Lozzo, Vida Dujmović, Fabrizio Frati, Tamara Mchedlidze, and Vincenzo Roselli. Drawing planar graphs with many collinear vertices. *Journal of Computational Geometry*, 9(1):94–130, 2018.

13 Vida Dujmović. The Utility of Untangling. *J. Graph Algorithms Appl.*, 21(1):121–134, 2017. doi:10.7155/jgaa.00407.

14 Vida Dujmović, Fabrizio Frati, Daniel Gonçalves, Pat Morin, and Günter Rote. Every Collinear Set is Free. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA19)*, 2019.

15 William Evans, Vincent Kusters, Maria Saumell, and Bettina Speckmann. Column Planarity and Partial Simultaneous Geometric Embedding. In Christian A. Duncan and Antonios Symvonis, editors, *Graph Drawing - 22nd International Symposium, GD 2014, Würzburg, Germany, September 24-26, 2014, Revised Selected Papers*, volume 8871 of *Lecture Notes in Computer Science*, pages 259–271. Springer, 2014. doi:10.1007/978-3-662-45803-7_22.

16 Emilio Di Giacomo, Giuseppe Liotta, and Tamara Mchedlidze. How many vertex locations can be arbitrarily chosen when drawing planar graphs? *CoRR*, abs/1212.0804, 2012. arXiv:1212.0804.

17 Xavier Goaoc, Jan Kratochvíl, Yoshio Okamoto, Chan-Su Shin, Andreas Spillner, and Alexander Wolff. Untangling a Planar Graph. *Discrete & Computational Geometry*, 42(4):542–569, 2009. `doi:10.1007/s00454-008-9130-6`.

18 Branko Grünbaum and Hansjoachim Walther. Shortness Exponents of Families of Graphs. *J. Comb. Theory, Ser. A*, 14(3):364–385, 1973. `doi:10.1016/0097-3165(73)90012-5`.

19 Bill Jackson. Longest cycles in 3-connected cubic graphs. *J. Combin. Theory Ser. B*, 41:17–26, 1986.

20 Mihyun Kang, Oleg Pikhurko, Alexander Ravsky, Mathias Schacht, and Oleg Verbitsky. Untangling planar graphs from a specified vertex position - Hard cases. *Discrete Applied Mathematics*, 159(8):789–799, 2011. `doi:10.1016/j.dam.2011.01.011`.

21 Goos Kant and Hans L. Bodlaender. Triangulating Planar Graphs while Minimizing the Maximum Degree. *Inf. Comput.*, 135(1):1–14, 1997. `doi:10.1006/inco.1997.2635`.

22 Qinghai Liu, Xingxing Yu, and Zhao Zhang. Circumference of 3-connected cubic graphs. *J. Comb. Theory, Ser. B*, 128:134–159, 2018. `doi:10.1016/j.jctb.2017.08.008`.

23 János Pach and Gábor Tardos. Untangling a Polygon. *Discrete & Computational Geometry*, 28(4):585–592, 2002. `doi:10.1007/s00454-002-2889-y`.

24 Alexander Ravsky and Oleg Verbitsky. On collinear sets in straight line drawings. *CoRR*, abs/0806.0253, 2008. `arXiv:0806.0253`.

25 Alexander Ravsky and Oleg Verbitsky. On Collinear Sets in Straight-Line Drawings. In Petr Kolman and Jan Kratochvíl, editors, *37th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2011)*, volume 6986 of *LNCS*, pages 295–306. Springer, 2011. `doi:10.1007/978-3-642-25870-1_27`.

26 Neil Robertson, Daniel P. Sanders, Paul D. Seymour, and Robin Thomas. The Four-Colour Theorem. *J. Comb. Theory, Ser. B*, 70(1):2–44, 1997. `doi:10.1006/jctb.1997.1750`.

27 Peter Guthrie Tait. Remarks on the colouring of maps. *Proc. Roy. Soc. Edinburgh Sect. A*, 10:729, 1880.

28 W. T. Tutte. On Hamilton circuits. *J. Lond. Math. Soc.*, 21:98–101, 1946.

29 Mamoru Watanabe. Open problem. 5th Czech–Slovak Symposium on Combinatorics, 1998.