Lower Bounds for Electrical Reduction on Surfaces*

Hsien-Chih Chang† Marcos Cossarini‡ Jeff Erickson§

March 22, 2019

Abstract

We strengthen the connections between *electrical transformations* and *homotopy* from the planar setting—observed and studied since Steinitz—to arbitrary surfaces with punctures. As a result, we improve our earlier lower bound on the number of electrical transformations required to reduce an n-vertex graph on surface in the worst case [SOCG 2016] in two different directions. Our previous $\Omega(n^{3/2})$ lower bound applies only to *facial* electrical transformations on plane graphs with *no terminals*. First we provide a stronger $\Omega(n^2)$ lower bound when the planar graph has two or more terminals, which follows from a quadratic lower bound on the number of homotopy moves in the annulus. Our second result extends our earlier $\Omega(n^{3/2})$ lower bound to the wider class of *planar* electrical transformations, which preserve the planarity of the graph but may delete cycles that are not faces of the given embedding. This new lower bound follows from the observation that the *defect* of the medial graph of a planar graph is the same for all its planar embeddings.

*This work was partially supported by NSF grant CCF-1408763. We also greatly appreciate the support from Labex Bezout. The conference version of the paper appears in the Proceedings of the 35th International Symposium on Computational Geometry (SoCG 2019).

†Department of Computer Science, Duke University, USA. This work was initiated when the author was affiliated with University of Illinois at Urbana-Champaign.

‡Laboratoire d'Analyse et de Mathématiques Appliquées, Université Paris-Est Marne-la-Vallée, France. This work was initiated when the author was affiliated with Instituto de Matemática Pura e Aplicada, Brazil.

§Department of Computer Science, University of Illinois at Urbana-Champaign, USA.

1 Introduction

Consider the following set of local operations performed on any graph:

- *Leaf contraction*: Contract the edge incident to a vertex of degree 1.
- *Loop deletion*: Delete the edge of a loop.
- *Series reduction*: Contract either edge incident to a vertex of degree 2.
- *Parallel reduction*: Delete one of a pair of parallel edges.
- $Y\to\Delta$ *transformation*: Delete a degree-3 vertex and connect its neighbors with three new edges.
- $\Delta\to Y$ *transformation*: Delete edges of a 3-cycle and join its vertices to a new vertex.

These operations and their inverses, which we call ***electrical transformations*** following Colin de Verdière *et al.* [13], have been used for over a century to analyze electrical networks [31]. Steinitz [45, 46] proved that any planar network can be reduced to a single vertex using these operations. Several decades later, Epifanov [17] proved that any planar graph with two special vertices called *terminals* can be similarly reduced to a single edge between the terminals; simpler algorithmic proofs of Epifanov's theorem were later given by Feo [19], Truemper [50, 51], and Feo and Provan [20]. These results have since been extended to planar graphs with more than two terminals [3, 14, 21, 22] and to some families of non-planar graphs [21, 52]. See Chang's thesis [8] for a history of the problem.

Despite decades of prior work, the complexity of the reduction process is still poorly understood. Steinitz's proof implies that $O(n^2)$ electrical transformations suffice to reduce any n-vertex planar graph to a single vertex; Feo and Provan's algorithm reduces any 2-terminal planar graph to a single edge in $O(n^2)$ steps. While these are the best upper bounds known, several authors have conjectured that they can be improved [3, 20, 21]. Without any restrictions on which transformations are permitted, the only known lower bound is the trivial $\Omega(n)$. However, Chang and Erickson recently proved that if all transformations are required to be *facial*, meaning any deleted cycle must be a face of the given embedding, then reducing a plane graph without terminals to a single vertex requires $\Omega(n^{3/2})$ steps in the worst case [10]. This is obtained by studying the relation between facial electrical transformations and *homotopy moves*, a set of operations performed on the medial graph of the input.

In this paper, we extend our earlier lower bound for electrical transformations in two directions. To this end, first we study multicurves on surfaces under electrical and homotopy moves; multicuves are in one-to-one correspondence with medial graphs of graph embeddings. Specifically, in Section 3 we prove that the set of *tight* multicurves under electrical moves and under homotopy moves is identical. As a consequence, any surface-embedded graph can be reduced without ever increasing its number of edges. Previously such property is only known to hold for plane graphs [10, 35].

Next, we consider plane graphs with two terminals. In this setting, leaf deletions, series reductions, and $Y\to\Delta$ transformations that delete terminals are forbidden. We prove in Section 4 that $\Omega(n^2)$ facial electrical transformations are required in the worst case to reduce a 2-terminal plane graph *as much as possible*. Not every 2-terminal plane graph can be reduced to a single edge between the terminals using only facial electrical transformations. However, we show that any 2-terminal plane graph can be reduced to a unique minimal graph called a *bullseye* using a finite number of facial electrical transformations. Our lower bound ultimately relies on a recent $\Omega(n^2)$ lower bound on the number of homotopy moves required to tighten a contractible closed curve in the annulus [12].

In Section 5, we consider a wider class of electrical transformations that preserve the planarity of the graph, but are not necessarily facial. Our second main result is that $\Omega(n^{3/2})$ *planar* electrical transformations are required to reduce a planar graph (without terminals) to a single vertex in the worst case. Like our earlier lower bound for *facial* electrical transformations, our proof ultimately reduces to the study of a certain curve invariant, called the *defect*, of the medial graph of a given *unicursal* plane graph G. A key step in our new proof is the following surprising observation: Although the definition of the medial graph of G depends on the embedding of G, the defect of the medial graph is the same for all planar embeddings of G.

2 Background

2.1 Types of electrical transformations

We distinguish between three increasingly general types of electrical transformations in plane graphs: *facial*, *crossing-free*, and *arbitrary*. (For ease of presentation, we assume throughout the paper that plane graphs are actually embedded on the *sphere* instead of the plane.)

An electrical transformation in a graph G embedded on a surface Σ is ***facial*** if any deleted cycle is a face of G. All leaf contractions, series reductions, and $Y\to\Delta$ transformations are facial, but loop deletions, parallel reductions, and $\Delta\to Y$ transformations may not be facial. Facial electrical transformations form three dual pairs, as shown in Figure 2.1; for example, any series reduction in G is equivalent to a parallel reduction in the dual graph G^*.

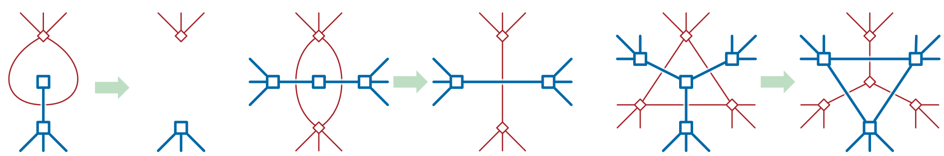

Figure 2.1. Facial electrical transformations in a plane graph G and its dual G^*.

An electrical transformation in G is ***crossing-free*** if it preserves the embeddability of the underlying graph into the same surface. Equivalently, an electrical transformation is crossing-free if the vertices of the cycle deleted by the transformation are all incident to a common face of G. All facial electrical transformations are trivially crossing-free, as are all loop deletions and parallel reductions. If the graph embeds in the plane, crossing-free electrical transformations are also called ***planar***. The only non-crossing-free electrical transformation is a $\Delta\to Y$ transformation whose three vertices are *not* incident to a common face; any such transformation introduces a $K_{3,3}$-minor into the graph, connecting the three vertices of the Δ to an interior vertex, an exterior vertex, and the new Y vertex.

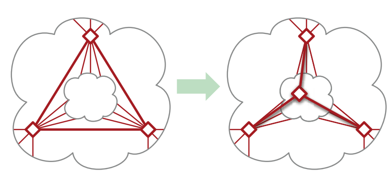

Figure 2.2. A non-planar $\Delta\to Y$ transformation.

2.2 Multicurves and medial graphs

A ***surface*** is a 2-manifold with or without punctures. Formally, a ***closed curve*** in a surface Σ is a continuous map $\gamma\colon S^1 \to \Sigma$. A closed curve is ***simple*** if it is injective. A ***multicurve*** is a collection of one or more closed curves. We consider only *generic* multicurves, which are injective except at a finite number of (self-)intersections, each of which is a transverse double point. A multicurve is ***connected*** if its image in the surface is connected. The image of any (non-simple) multicurve has a natural structure as a 4-regular map, whose ***vertices*** are the self-intersection points of the curves, ***edges*** are maximal subpaths between vertices, and ***faces*** are components of the complement of the curves in the surface. We do not distinguish between multicurves whose images are combinatorially equivalent maps.

The ***medial graph*** $G^\times$ of an embedded graph G is another embedded graph whose vertices correspond to the edges of G, and two vertices of $G^\times$ are connected by an edge if the corresponding edges in G are consecutive in cyclic order around some vertex, or equivalently, around some face in G. Every vertex in every medial graph has degree 4; thus, every medial graph is the image of a multicurve. Conversely, image of a non-simple multicuvre is the medial graph of some surface-embedded graph if the faces of the multicurve can be two-colored; in particular, when the surface is a sphere, the image of every non-simple multicurve is the medial graph of some plane graph. We call an embedded graph G ***unicursal*** if its medial graph $G^\times$ is the image of a single closed curve.

Smoothing a multicurve γ at a vertex x replaces the intersection of γ with a small neighborhood of x with two disjoint simple paths, so that the result is another 4-regular embedded graph. There are two possible smoothings at each vertex. More generally, a ***smoothing*** of γ is any multicurve obtained by smoothing a subset of its vertices. For any embedded graph G, the smoothings of the medial graph $G^\times$ are precisely the medial graphs of minors of G.

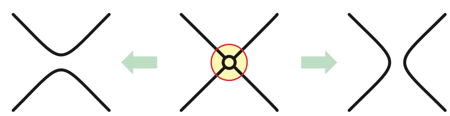

Figure 2.3. Two possible smoothings of a vertex.

2.3 Local moves

A ***homotopy*** between two curves γ and γ' on the same surface Σ is a continuous deformation from one curve to the other, formally defined as a continuous function $H : S^1 \times [0,1] \to \Sigma$ such that $H(\cdot, 0) = \gamma$ and $H(\cdot, 1) = \gamma'$. The definition of homotopy extends naturally to multicurves. Classical topological arguments imply that two multicurves are homotopic if and only if one can be transformed into the other by a finite sequence of ***homotopy moves*** (shown in Figure 2.4). Notice that a 1→0 move is applied to an empty ***loop***, and a 2→0 move is applied on an empty ***bigon***. A multicurve is ***homotopically tight*** (or ***h-tight*** for short) if no sequence of homotopy moves leads to a multicurve with fewer vertices.

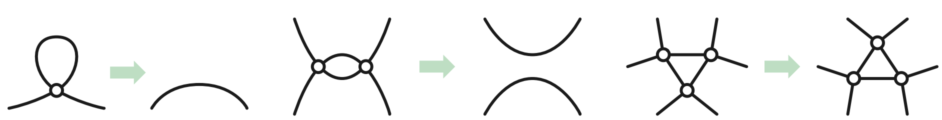

Figure 2.4. Homotopy moves 1→0, 2→0, and 3→3.

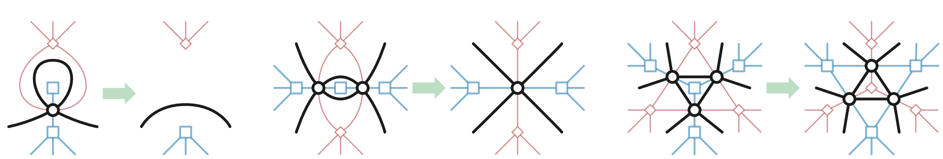

Figure 2.5. Electrical moves 1→0, 2→1, and 3→3.

Facial electrical transformations in any embedded graph G correspond to local operations in the medial graph $G^\times$ that closely resemble homotopy moves. We call these **1→0**, **2→1**, and **3→3** moves, where the numbers before and after each arrow indicate the number of local vertices before and after the move. We collectively refer to these operations and their inverses as ***electrical moves***. A multicurve is ***electrically tight*** (or ***e-tight*** for short) if no sequence of electrical moves leads to another multicurve with fewer vertices. For multicurves on surfaces with boundary, both homotopy moves and electrical moves performed on boundary faces are forbidden. The fact that we use same name *tight* for both homotopy moves and electrical moves is not a coincidence; we will justify its usage in Section 3.2.

3 Connection between electrical and homotopy moves

For any connected multicurve (or 4-regular embedded graph) γ on surface Σ,

- let $\boldsymbol{X(\gamma)}$ denote the minimum number of electrical moves required to tighten γ,
- let $\boldsymbol{H^{\downarrow}(\gamma)}$ denote the minimum number of homotopy moves required to tighten γ, without ever increase the number of vertices; that is, no $0{\to}1$ and $0{\to}2$ moves are allowed.
- let $\boldsymbol{H(\gamma)}$ denote the minimum number of homotopy moves required to tighten γ.

It is not immediately obvious whether a multicurve γ that is tight under monotonic homotopy moves could be further tightened by allowing $0{\to}1$ and $0{\to}2$ moves or not. Hass and Scott [28] and de Graaf and Schrijver [26] independently proved that any multicurve γ can be tightened using monotonic homotopy moves, which implies that $H^{\downarrow}(\gamma) = 0$ if and only if $H(\gamma) = 0$. In other words, (standard) homotopy moves and monotonic homotopy moves share the same set of tight multicurves. Now $H^{\downarrow}(\gamma) \geq H(\gamma)$ follows for any multicurve γ.

3.1 Smoothing lemma

We would like to compare $X(\gamma)$ with $H^{\downarrow}(\gamma)$ and $H(\gamma)$. The following key lemma follows from close reading of proofs by Truemper [50, Lemma 4] and several others [3, 21, 33, 35] that every minor of a ΔY-reducible graph is also ΔY-reducible. A proof to some special cases at the level of medial curves can be found in de Graaf [23, Proposition 5.1]. For the sake of completeness, we include a proof in Appendix B.

Lemma 3.1 (Chang and Erickson [10, Lemma 3.1]). *Let γ be any connected multicurve on surface Σ, and let $\check{\gamma}$ be a connected smoothing of γ. Applying any sequence of N electrical moves to γ to obtain γ'. Then one can apply a similar sequence of electrical moves of length at most N to $\check{\gamma}$ to obtain a (possibly trivial) connected smoothing $\check{\gamma}'$ of γ'.*

As a remark, using similar argument one can recover a result by Newmann-Coto [34]: any homotopy from multicurve γ to another multicurve γ' that never removes vertices can be turned into a homotopy from a smoothing of γ to a smoothing of γ'. Chambers and Liokumovich [7] studied a similar problem where one wants to convert a homotopy between two *simple* curves on surface into an *isotopy*, without increasing the length of any intermediate curve by too much. They showed that the desired isotopy can be obtained from a clever Euler-tour argument on the graph of all possible complete smoothings of the intermediate curves.

Using Lemma 3.1 one can show that $X(\gamma) \geq H^{\downarrow}(\gamma)$ for every planar curve γ, a result implicit in the work of Noble and Welsh [35] and formally proved by Chang and Erickson [10].

Lemma 3.2 (Smoothing Lemma [10]). *$X(\check{\gamma}) \leq X(\gamma)$ for every connected smoothing $\check{\gamma}$ of every connected multicurve γ in the plane.*

Lemma 3.3 (Monotonicity Lemma [10]). *For every connected multicurve γ, there is a minimum-length sequence of electrical moves that simplifies γ to a simple closed curve that does not contain $0{\to}1$ or $1{\to}2$ moves.*

Lemma 3.4 (Electrical-Homotopy Inequality [10]). $X(\gamma) \geq H^{\downarrow}(\gamma)$ *for every planar curve γ.*

3.2 Equivalence of tightness

One of the main obstacles to generalize Lemmas 3.2, 3.3, and 3.4 to curves on arbitrary surface is that again we do not know *a priori* whether the set of tight multicurves under electrical moves is the same as those under homotopy moves. Such problem did not exist in the planar setting as all planar multicurves can be tightened to simple curves using either electrical or homotopy moves. We first show that every electrically tight multicurve is also homotopically tight.

Lemma 3.5. *Let γ be a connected multicurve on an arbitrary surface Σ. If γ is electrically tight, then γ is homotopically tight.*

Proof: Let γ be a connected multicurve in some arbitrary surface, and suppose γ is not homotopically tight. Results of Hass and Scott [28] and de Graaf and Schrijver [26] imply that γ can be tightened by a finite sequence of homotopy moves that never increases the number of vertices. In particular, applying some finite sequence of $3{\to}3$ moves to γ creates either an empty loop, which can be removed by a $1{\to}0$ move, or an empty bigon, which can be removed by either a $2{\to}0$ move or a $2{\to}1$ move. Thus, γ is not electrically tight. $\square$

However, for the reverse direction, we don't have a similar monotonicity result for electrical moves on arbitrary surfaces. A careful reading of the sequence of work by de Graaf and Schrijver [24, 25, 26, 39, 40, 41, 42] leads to a five-way equivalence that shows the two versions of tightness coincide when the given curve is *primitive*. Unfortunately their results do not generalize as some of the equivalences break down with the presence of non-primitive counterexamples. See Appendix A for more details.

Routing set. Inspired by the routing problem studied by de Graaf and Schrijver [25], we introduce the notion of *routing set*. Despite its naïve look, the routing set satisfies a crucial property that encapsulates the whole difficulty of the problem, which allows us to bypass the heavy machinery developed for the primitive case. We then use the established equivalence of tightness to derive the monotonicity lemma for electrical moves on arbitrary multicurves.

For any multicurve γ, the ***routing set*** of γ is the following collection of homotopy classes:

$$\textbf{\textit{route}}(\boldsymbol{\gamma}) \coloneqq \left\{ [\check{\gamma}] \mid \check{\gamma} \text{ is a smoothing of } \gamma \right\}.$$

Each homotopy class in *route*(γ) is referred as a ***route*** of γ.

Lemma 3.6. *Routing set of γ is invariant under electrical moves for any multicurve γ.*

Proof: Let γ' be the multicurve obtained from performing one electrical move to γ. Because electrical moves are closed under inverses, we only need to prove that *route*(γ) $\subseteq$ *route*(γ').

Let $\check{\gamma}$ be an arbitrary smoothing of γ; $[\check{\gamma}]$ is in *route*(γ) by definition. By Lemma 3.1, one can obtain a smoothing $\check{\gamma}'$ of γ' that is at most one electrical move away from $\check{\gamma}$.[1] In particular, $[\check{\gamma}']$ is in *route*(γ').

[1] Although Lemma 3.1 is stated with respect to *connected* smoothings, the proof of the lemma (see Appendix B) reveals that similar statement holds for arbitrary smoothings by allowing an additional $0{\to}0$ move that creates/contracts simple cycles. In particular, such move does not change the homotopy class of a multicurve.

If $\check{\gamma}'$ is equal to $\check{\gamma}$ or is obtained from $\check{\gamma}$ using a $1{\to}0$, $0{\to}1$, or $3{\to}3$ move, then immediately we have $[\check{\gamma}] = [\check{\gamma}']$ to be a route in $route(\gamma')$. If $\check{\gamma}'$ is obtained from $\check{\gamma}$ using a $2{\to}1$ move, consider the multicurve $\check{\gamma}^\circ$ obtained from $\check{\gamma}$ by performing a $2{\to}0$ move (on the same empty bigon) instead. $\check{\gamma}^\circ$ is a smoothing of $\check{\gamma}'$, which in turn is a smoothing of γ'. Because $2{\to}0$ is a homotopy move, $[\check{\gamma}] = [\check{\gamma}^\circ]$ is a route in $route(\gamma')$. Similarly when $\check{\gamma}'$ is obtained from $\check{\gamma}$ using a $1{\to}2$ move, we consider $\check{\gamma}$ as a smoothing of $\check{\gamma}'$ thus $[\check{\gamma}]$ is a route in $route(\gamma')$. This concludes the proof. □

The ***intersection number*** of a homotopy class $[\gamma]$ is defined to be the minimum number of vertices among all curves homotopic to γ. The ***main routes*** of γ are those routes of γ that achieve the maximum intersection number.

Lemma 3.7. *Any homotopically tight multicurve is also electrically tight.*

Proof: Assume for contradiction that there is an h-tight multicurve γ that is not e-tight. Tighten γ using electrical moves to an e-tight multicurve γ' with less number of vertices than γ. Now by Lemma 3.6 the routing set of γ and γ' is the same; in particular, $[\gamma']$ is a main route of both γ and γ'. However since both γ and γ' are h-tight, the intersection number of $[\gamma]$ is strictly greater than the intersection number of $[\gamma']$ and thus $[\gamma']$ cannot be a main route of γ, a contradiction. □

3.3 Monotonicity of electrical moves

As a corollary of Lemma 3.7, we are ready to generalize the monotonicity lemma (Lemma 3.3) to multicurves on general surfaces.

Lemma 3.8. *Let γ be any connected multicurve γ on surface Σ, and let $\check{\gamma}$ be a connected smoothing of γ, satisfying $route(\gamma) = route(\check{\gamma})$. Then $X(\check{\gamma}) \le X(\gamma)$ holds.*

Proof: Let γ be a connected multicurve with $n(\gamma)$ vertices, and let $\check{\gamma}$ be a connected smoothing of γ. If $X(\gamma)$ equals to zero, then γ is both e-tight and h-tight by Lemma 3.5. The fact that $route(\gamma) = route(\check{\gamma})$ implies that $[\gamma]$ is a route of $\check{\gamma}$ and its intersection number is equal to $n(\gamma)$. If $\check{\gamma}$ is a proper smoothing of γ, then the intersection number of any route of $\check{\gamma}$ is strictly less then $n(\gamma)$, a contradiction. As a result, the only smoothing of γ satisfying the condition is γ itself, and therefore the inequality trivially holds.

Otherwise, applying a minimum-length sequence of electrical moves that tightens γ. By Lemma 3.1 there is another sequence of electrical moves of length at most $X(\gamma)$ that tightens $\check{\gamma}$. We immediately have $X(\check{\gamma}) \le X(\gamma)$ and the lemma is proved. □

Lemma 3.9. *For any connected multicurve γ, there is a minimum-length sequence of electrical moves that tightens γ that does not contain $0{\to}1$ or $1{\to}2$ moves.*

The proof follows almost verbatim from Lemma 3.3 after substituting Lemma 3.8 for Lemma 3.2 and applying Lemma 3.6.

Proof: Consider a minimum-length sequence of electrical moves that tights γ. For any integer $i \ge 0$, let γ_i denote the result of the first i moves in this sequence. Minimality of the tightening sequence implies that $X(\gamma_i)$ decreases as i grows. Now let i be an arbitrary index such that γ_i is obtained from performing a $0{\to}1$ or $1{\to}2$ move on γ_{i-1}. Then γ_{i-1} is a connected proper smoothing of γ_i, and by Lemma 3.6, $route(\gamma_{i-1}) = route(\gamma_i)$ holds. Now Lemma 3.8 implies that $X(\gamma_{i-1}) \le X(\gamma_i)$, a contradiction. □

4 Two-terminal plane graphs

Most applications of electrical reductions, starting with Kennelly's computation of effective resistance [31], designate two vertices of the input graph as *terminals* and require a reduction to a single edge between those terminals. In this context, electrical transformations that delete either of the terminals are forbidden; specifically: leaf contractions when the leaf is a terminal, series reductions when the degree-2 vertex is a terminal, and $Y \to \Delta$ transformations when the degree-3 vertex is a terminal. An important subtlety here is that not every 2-terminal planar graph can be reduced to a single edge using only *facial* electrical transformations. The simplest bad example is the three-vertex graph shown in Figure 4.1.

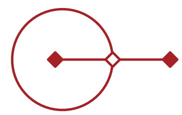

Figure 4.1. A facially irreducible 2-terminal plane graph; solid vertices are the terminals.

In this section, we show that in the worst case, $\Omega(n^2)$ facial electrical transformations are required to reduce a 2-terminal plane graph with n vertices *as much as possible*. The medial graph $G^\times$ of any 2-terminal plane graph G is properly considered as a multicurve embedded in the annulus; the faces of $G^\times$ that correspond to the terminals are removed from the surface. The main strategy is to lower bound $X(G^\times)$ by some function of $H(G^\times)$, then defer to the quadratic lower bound for untangling annular curve using homotopy moves [12]. To this end, we generalize Lemma 3.4 to annular curves; such result is obtained by the understanding of tight multicurves on the annulus.

First, we prove in Section 4.1 that any annular curve can be tightened to a unique family of curves. Next in Section 4.2, we generalize the results by Chang and Erickson [10], in particular the electrical-homotopy inequality (Lemma 3.4), to the annular case. We prove our quadratic lower bound in Section 4.3. Existing algorithms for reducing an arbitrary 2-terminal plane graphs to a single edge rely on an additional operation which we call a *terminal-leaf contraction*, in addition to facial electrical transformations. We discuss this subtlety in more detail in Section 4.4.

4.1 Tight annular curves

The ***winding number*** of a directed closed curve γ in the annulus is the number of times any generic path π from one (fixed) boundary component to the other crosses γ from left to right, minus the number of times π crosses γ from right to left. Two directed closed curves in the annulus are homotopic if and only if their winding numbers are equal.

The ***depth*** of any multicurve γ in the annulus is the minimum number of times a path from one boundary to the other crosses γ; thus, depth is essentially an unsigned version of winding number. Just as the winding number around the boundaries is a complete homotopy invariant for curves in the annulus, the depth turns out to be a complete invariant for electrical moves on the annular multicurves.

Lemma 4.1. *Electrical moves do not change the depth of any annular multicurve.*

For any integer $d > 0$, let $\boldsymbol{\alpha_d}$ denote the unique closed curve in the annulus with $d-1$ vertices and winding number d. Up to isotopy, this curve can be parametrized in the plane as

$$\alpha_d(\theta) := \Big((\cos(\theta)+2)\cos(d\theta),\ (\cos(\theta)+2)\sin(d\theta)\Big).$$

In the notation of our other papers [10, 11], α_d is the *flat torus knot* $T(d,1)$.

The following lemmas are direct consequences of Lemma 3.7; here we provide simple proofs using only winding number and depth of annular curves.

Lemma 4.2. *For any integer $d > 0$, the curve α_d is both h-tight and e-tight.*

Proof: Every connected multicurve in the annulus with either winding number d or depth d has at least $d+1$ faces (including the faces containing the boundaries of the annulus) and therefore, by Euler's formula, has at least $d-1$ vertices. □

Lemma 4.3. *If γ is an h-tight connected annular multicurve, then $\gamma = \alpha_d$ for some d.*

Proof: A multicurve in the annulus is h-tight if and only if its constituent curves are h-tight *and disjoint*. Thus, any *connected* h-tight multicurve is actually a single closed curve. Any two curves in the annulus with the same winding number are homotopic [30]. Finally, up to isotopy, α_d is the only closed curve in the annulus with winding number d and $d-1$ vertices [27, Lemma 1.12]. □

Corollary 4.4. *A connected multicurve γ in the annulus is e-tight if and only if $\gamma = \alpha_{depth(\gamma)}$; therefore, any annular multicurve γ is e-tight if and only if γ is h-tight.*

4.2 Smoothing lemma in the annulus

Equipped with the understanding of tight annular curves, we are ready to extend the results in Section 3.1 to the annulus.

Lemma 4.5. *For any connected smoothing $\check{\gamma}$ of any connected multicurve γ in the annulus, we have $X(\check{\gamma}) + \frac{1}{2}\,depth(\check{\gamma}) \le X(\gamma) + \frac{1}{2}\,depth(\gamma)$.*

Proof: Let γ be an arbitrary connected multicurve in the annulus, and let $\check{\gamma}$ be an arbitrary connected smoothing of γ. Without loss of generality, we can assume that γ is non-simple, since otherwise the lemma is vacuous.

If γ is already e-tight, then $\gamma = \alpha_d$ for some integer $d \ge 2$ by Corollary 4.4. (The curves α_0 and α_1 are simple.) First, suppose $\check{\gamma}$ is a connected smoothing of γ obtained by smoothing a single vertex x. The smoothed curve $\check{\gamma}$ contains a single empty loop if x is the innermost or outermost vertex of γ, or a single empty bigon otherwise. Applying one $1{\to}0$ or $2{\to}0$ move transforms $\check{\gamma}$ into the curve α_{d-2}, which is e-tight by Lemma 4.2. Thus we have $X(\check{\gamma}) = 1$ and $depth(\check{\gamma}) = d-2$, which implies $X(\check{\gamma}) + \frac{1}{2}\,depth(\check{\gamma}) = X(\gamma) + \frac{1}{2}\,depth(\gamma)$. As for the general case when $\check{\gamma}$ is obtained from γ by smoothing more than one vertices, the statement follows from the previous case by induction on the number of smoothed vertices.

If γ is not e-tight, applying a minimum-length sequence of electrical moves that tightens γ into some curve γ'. By Lemma 3.1 there is another sequence of electrical moves of length at most $X(\gamma)$ that tightens $\check{\gamma}$ to some connected smoothing $\check{\gamma}'$ of γ', which can be further tightened electrically to an e-tight curve using arguments in the previous paragraph because γ' is e-tight. This implies that $X(\check{\gamma}) \le X(\gamma) + \frac{1}{2}(depth(\gamma') - depth(\check{\gamma}'))$. By Lemma 4.1, γ and γ' have the same depth, and $\check{\gamma}$ and $\check{\gamma}'$ have the same depth. Therefore $X(\check{\gamma}) + \frac{1}{2}\,depth(\check{\gamma}) \le X(\gamma) + \frac{1}{2}\,depth(\gamma)$ and the lemma is proved. □

Lemma 4.6. *For every connected multicurve γ in the annulus, there is a minimum-length sequence of electrical moves that tightens γ to $\alpha_{depth(\gamma)}$ without $0{\to}1$ or $1{\to}2$ moves.*

The proof follows almost verbatim from Lemma 3.3 and 3.9 after substituting Lemma 4.5 for Lemma 3.2.

Proof: Consider a minimum-length sequence of electrical moves that tightens an arbitrary connected multicurve γ in the annulus. For any integer $i \geq 0$, let γ_i denote the result of the first i moves in this sequence. Suppose γ_i has one more vertex than γ_{i-1} for some index i. Then γ_{i-1} is a connected proper smoothing of γ_i, and $depth(\gamma_i) = depth(\gamma_{i-1})$ by Lemma 4.1; so Lemma 4.5 implies that $X(\gamma_{i-1}) \leq X(\gamma_i)$, contradicting our assumption that the reduction sequence has minimum length. □

Lemma 4.7. *$X(\gamma) + \frac{1}{2}\,depth(\gamma) \geq H^{\downarrow}(\gamma) \geq H(\gamma)$ for every closed curve γ in the annulus.*

Proof: Let γ be a closed curve in the annulus. If γ is already e-tight, then $X(\gamma) = H^{\downarrow}(\gamma) = 0$ by Lemma 3.5 (or Corollary 4.4), so the lemma is trivial. Otherwise, consider a minimum-length sequence of electrical moves that tightens γ. By Lemma 4.6, we can assume that the first move in the sequence is neither 0→1 nor 1→2. If the first move is 1→0 or 3→3, the theorem immediately follows by induction on $X(\gamma)$, since by Lemma 4.1 neither of these moves changes the depth of the curve.

The only interesting first move is 2→1. Let γ' be the result of this 2→1 move, and let γ° be the result if we perform the 2→0 move on the same empty bigon instead. The minimality of the sequence implies $X(\gamma) = X(\gamma') + 1$, and we trivially have $H^{\downarrow}(\gamma) \leq H^{\downarrow}(\gamma^\circ) + 1$. Because γ is a single curve, γ° is also a single curve and therefore a connected proper smoothing of γ'. Thus, Lemma 4.1, Lemma 4.5, and induction on the number of vertices imply

$$\begin{aligned} X(\gamma) + \frac{1}{2}\,depth(\gamma) &= X(\gamma') + \frac{1}{2}\,depth(\gamma') + 1 \\ &\geq X(\gamma^\circ) + \frac{1}{2}\,depth(\gamma^\circ) + 1 \\ &\geq H^{\downarrow}(\gamma^\circ) + 1 \\ &\geq H^{\downarrow}(\gamma), \end{aligned}$$

which completes the proof. □

4.3 Quadratic lower bound

Bullseyes. For any $k > 0$, let $\boldsymbol{B_k}$ denote the 2-terminal plane graph that consists of a path of length k between the terminals, with a loop attached to each of the $k-1$ interior vertices, embedded so that collectively they form concentric circles that separate the terminals. We call each graph B_k a ***bullseye***. For example, B_1 is just a single edge; B_2 is shown in Figure 4.1; and B_4 is shown on the left in Figure 4.2. The medial graph $B_k^\times$ of the kth bullseye is the curve α_{2k}. Because different bullseyes have different medial depths, Lemma 4.1 implies that no bullseye can be transformed into any other bullseye by facial electrical transformations.

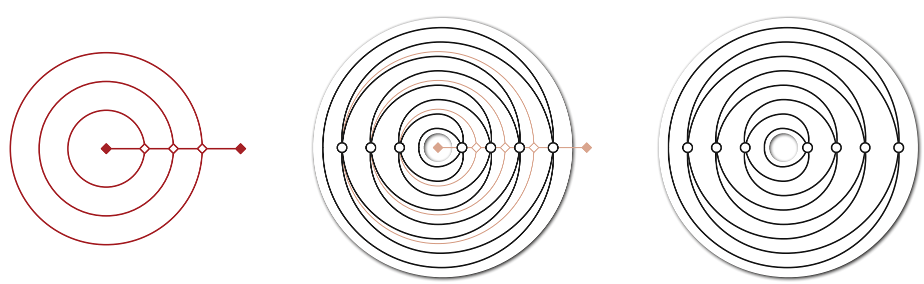

Figure 4.2. The bullseye graph B_4 and its medial graph α_8.

The following corollary is now immediate from the electrical-homotopy inequality for annular curves (Lemma 4.7).

Theorem 4.8. *Let G be a 2-terminal plane graph, and let γ be any unicursal smoothing of $G^\times$. Reducing G to a bullseye requires at least $H(\gamma) - \frac{1}{2}\,depth(\gamma)$ facial electrical transformations.*

Chang *et al.* [12] presented an infinite family of contractible curves in the annulus parametrized by their number of vertices n that require $\Omega(n^2)$ homotopy moves to tighten. Every contractible curve is the medial graph of some 2-terminal plane graph (because they have even depth and thus the faces can be two-colored [47]). Euler's formula implies that every n-vertex curve in the annulus has exactly $n+2$ faces (including the boundary faces) and therefore has depth at most $n+1$.

Corollary 4.9. *Reducing a 2-terminal plane graph to a bullseye requires $\Omega(n^2)$ facial electrical transformations in the worst case.*

4.4 Terminal-leaf contractions

The electrical reduction algorithms of Feo [19], Truemper [50], and Feo and Provan [20] rely exclusively on facial electrical transformations, plus one additional operation.

- *Terminal-leaf contraction*: Contract the edge incident to a *terminal* vertex with degree 1. The neighbor of the deleted terminal becomes a new terminal.

Terminal-leaf contractions are also called *FP-assignments*, after Feo and Provan [14, 21, 22]. Later algorithms for reducing plane graphs with three or four terminals [3, 14, 22] also use only facial electrical transformations and terminal-leaf contractions.

Formally, terminal-leaf contractions are *not* electrical transformations, as they can change the value one wants to compute. For example, if the edges in the graph shown in Figure 4.1 represent 1Ω resistors, a terminal-leaf contraction changes the effective resistance between the terminals from 2Ω to 1Ω. However, both Gilter [21] and Feo and Provan [20] observed that any sequence of facial electrical transformations and terminal-leaf contractions can be simulated on the fly by a sequence of *planar* electrical transformations. Specifically, we simulate the first leaf contraction at either terminal by simply marking that terminal and proceeding as if its unique neighbor were a terminal. Later electrical transformations involving the neighbor of a marked terminal may no longer be facial, but they will still be planar; terminal-leaf contractions at the unique neighbor of a marked terminal become series reductions. At the end of the sequence of transformations, we perform a final series reduction at the unique neighbor of each marked terminal.

Unfortunately, terminal-leaf contractions change both the depth of the medial graph and the curve invariants that imply the quadratic homotopy lower bound. As a result, our quadratic lower bound proof breaks down if we allow terminal-leaf contractions.

5 Planar electrical transformations

Finally, we extend our earlier $\Omega(n^{3/2})$ lower bound for reducing plane graphs—*without* terminals using only facial electrical transformations—to the larger class of *planar* electrical transformations. Recall that a plane graph G *unicursal* if its medial graph $G^\times$ is the image of a single closed curve. As in our earlier work [10], we analyze electrical transformations in an unicursal plane graph G in terms of a certain invariant of the medial graph of G called ***defect***, first introduced by Aicardi [2] and Arnold [4, 5]. Our extension to non-facial electrical transformations is based on the following surprising observation: Although the medial graph of G depends on its embedding, the *defect* of the medial graph of G does not.

Theorem 5.1. *Let G and H be planar embeddings of the same abstract planar graph. If G is unicursal, then H is unicursal and $defect(G^\times) = defect(H^\times)$.*

The goal of the section is to prove Theorem 5.1.

5.1 Defect

Let γ be an arbitrary closed curve on the sphere. Choose an arbitrary basepoint $\gamma(0)$ and an arbitrary orientation for γ. For any vertex x of γ, we define $\operatorname{sgn}(x) = +1$ if the first traversal through x crosses the second traversal from right to left, and $\operatorname{sgn}(x) = -1$ otherwise. Two vertices x and y are ***interleaved***, denoted $\boldsymbol{x \between y}$, if they alternate in cyclic order—x, y, x, y—along γ. Finally, following Polyak [36], we can define

$$\mathit{defect}(\gamma) \;:=\; -2\sum_{x \between y} \operatorname{sgn}(x)\cdot\operatorname{sgn}(y),$$

where the sum is taken over all interleaved pairs of vertices of γ.

Trivially, every simple closed curve has defect zero. Straightforward case analysis [36] implies that the defect of a curve does not depend on the choice of basepoint or orientation. Moreover, any homotopy move changes the defect of a curve by at most 2; see the paper by Chang and Erickson [10, Section 2.1] for an explicit case breakdown. Defect is also preserved by any homeomorphism from the sphere to itself, including reflection.

5.2 Navigating between planar embeddings

Short history of planar embeddings. A classical result of Adkisson [1] and Whitney [53] is that every 3-connected planar graph has an essentially unique planar embedding. Mac Lane [32] described how to count the planar embeddings of any biconnected planar graph, by decomposing it into its triconnected components. Stallmann [43, 44] and Cai [6] extended Mac Lane's algorithm to arbitrary planar graphs, by decomposing them into biconnected components. Mac Lane's decomposition is also the basis of the SPQR-tree data structure of Di Battista and Tamassia [15, 16], which encodes all planar embeddings of an arbitrary planar graph.

Whitney [49, 54] showed that any planar embedding of a 2-connected planar graph G can be transformed into any other embedding by a finite sequence of ***split reflections***, defined as follows. A ***split curve*** is a simple closed curve σ whose intersection with the embedding of G consists of two vertices x and y; without loss of generality, σ is a circle with x and y at opposite points. A split reflection modifies the embedding of G by reflecting the subgraph inside σ across the line through x and y.

Lemma 5.2. *Let G be an arbitrary 2-connected planar graph. Any two planar embeddings of G can be transformed into one other by a finite sequence of split reflections.*

To navigate among the planar embeddings of *arbitrary* connected planar graphs, we need two additional operations. First, we allow split curves that intersect G at only a single cut vertex; a ***cut reflection*** modifies the embedding of G by reflects the subgraph inside such a curve. More interestingly, we also allow degenerate split curves that pass through a cut vertex x of G *twice*, but are otherwise simple and disjoint from G. The interior of a degenerate split curve σ is an open topological disk. A ***cut eversion*** is a degenerate split reflection that everts the embedding of the subgraph of G inside such a curve, intuitively by mapping the interior of σ to an open circular disk (with two copies of x on its boundary), reflecting the interior subgraph, and then mapping the resulting embedding back to the interior of σ. Structural results of Stallman [43, 44] and Di Battista and Tamassia [16, Section 7] imply the following.

Lemma 5.3. *Let G be an arbitrary connected planar graph. Any planar embedding of G can be transformed into any other planar embedding of G by a finite sequence of split reflections, cut reflections, and cut eversions.*

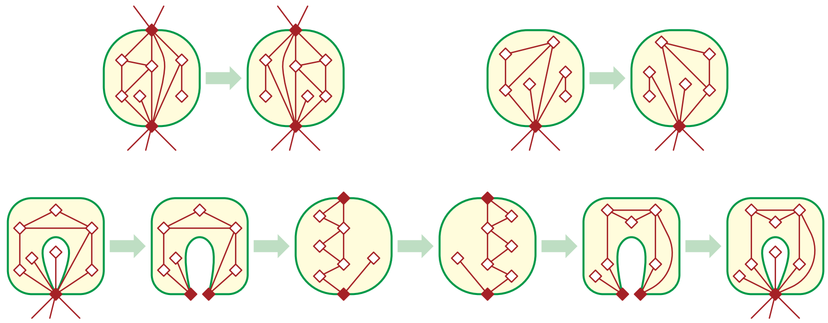

Figure 5.1. Top row: A regular split reflection and a cut reflection. Bottom row: a cut eversion.

5.3 Tangle flips

Now consider the effect of the operations stated in Lemma 5.3 on the medial graph $G^\times$. By assumption, G is unicursal so that $G^\times$ is a single closed curve. Let σ be any (possibly degenerate) split curve for G. Embed $G^\times$ so that every medial vertex lies on the corresponding edge in G, and every medial edge intersects σ at most once. By the Jordan curve theorem, we can assume without loss of generality that σ is a circle, and that the intersection points $\gamma \cap \sigma$ are evenly spaced around σ. A ***tangle*** of γ is the intersection of γ with either disk bounded by σ; each tangle consists of one or more subpaths of γ called ***strands***. We arbitrarily refer to the two tangles defined by σ as the *interior* and *exterior* tangles of σ. Split curve σ intersects at most four edges of $G^\times$, so the tangle of $G^\times$ inside σ has at most two strands. Moreover, reflecting (or everting) the subgraph of G inside σ induces a ***flip*** of this tangle of $G^\times$. Any tangle can be *flipped* by reflecting the disk containing it, so that each strand endpoint maps to a different strand endpoint; see Figure 5.2. Straightforward case analysis implies that flipping any tangle of $G^\times$ with at most two strands transforms $G^\times$ into another closed curve; see Figure 5.3.

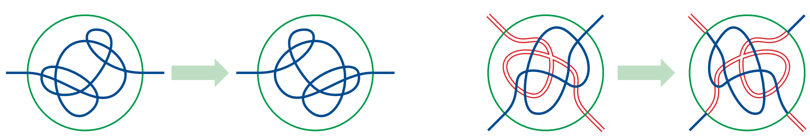

Figure 5.2. Flipping tangles with one and two strands.

Lemma 5.4. *Let γ be an arbitrary closed curve on the sphere. Flipping any tangle of γ with one strand yields another closed curve γ' with defect(γ') = defect(γ).*

Proof: Let σ be a simple closed curve that crosses γ at exactly two points. These points decompose σ into two subpaths $\alpha \cdot \beta$, where α is the unique strand of the interior tangle and β is the unique strand of the exterior tangle. Let Σ denote the interior disk of σ, and let $\phi : \Sigma \to \Sigma$ denote the homeomorphism that flips the interior tangle. Flipping the interior tangle yields the closed curve $\gamma' := rev(\phi(\alpha)) \cdot \beta$, where *rev* denotes path reversal.

No vertex of α is interleaved with a vertex of β; thus, two vertices in γ' are interleaved if and only if the corresponding vertices in γ are interleaved. Every vertex of $rev(\phi(\alpha))$ has the same sign as the corresponding vertex of α, since both the orientation of the vertex and the order of traversals through the vertex changed. Thus, every vertex of γ' has the same sign as the corresponding vertex of γ. We conclude that defect(γ') = defect(γ). □

A tangle is ***tight*** if each strand is simple and each pair of strands crosses at most once. Any tangle can be ***tightened***—that is, transformed into a tight tangle—by continuously deforming the strands without crossing σ or moving their endpoints, and therefore by a finite sequence of homotopy moves. Let $\boldsymbol{\gamma \Cap \sigma}$ and $\boldsymbol{\gamma \Cup \sigma}$ denote the closed curves that result from tightening the interior and exterior tangles of σ, respectively.[2] The following lemma that flipping any 2-strand tangle does not change its defect follows from our inclusion-exclusion formula for defect [9, Lemma 5.4]; we give a simpler proof here to keep the paper self-contained.

Lemma 5.5. *Let γ be an arbitrary closed curve on the sphere. Flipping any tangle of γ with two strands yields another closed curve γ' with defect(γ') = defect(γ).*

Proof: Let σ be a simple closed curve that crosses γ at exactly four points. These four points naturally decompose γ into four subpaths $\alpha \cdot \delta \cdot \beta \cdot \varepsilon$, where α and β are the strands of the interior tangle of σ, and δ and ε are the strands of the exterior tangle. Flipping the interior tangle either exchanges α and β, reverses α and β, or both; see Figure 5.3. In every case, the result is a single closed curve γ'. We classify each vertex of γ as *interior* if it lies on α and/or β, and *exterior* otherwise. Similarly, we classify pairs of interleaved vertices are either interior, exterior, or mixed.

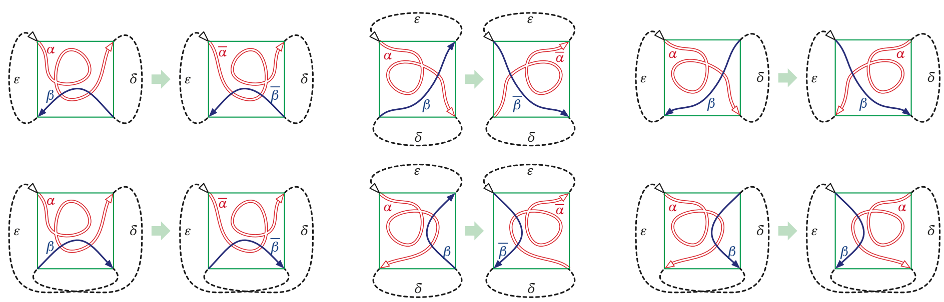

Figure 5.3. Flipping all six types of 2-strand tangle.

An interior vertex x and an exterior vertex y are interleaved if and only if x is an intersection point of α and β and y is an intersection point of δ and ε. Thus, the total contribution of mixed vertex pairs to Polyak's formula *defect*$(\gamma) = -2\sum_{x \between y} \operatorname{sgn}(x) \cdot \operatorname{sgn}(y)$ is

$$-2 \sum_{x \in \alpha \cap \beta} \sum_{y \in \delta \cap \varepsilon} \operatorname{sgn}(x) \cdot \operatorname{sgn}(y) \;=\; -2 \left(\sum_{x \in \alpha \cap \beta} \operatorname{sgn}(x) \right) \left(\sum_{y \in \delta \cap \varepsilon} \operatorname{sgn}(y) \right).$$

Consider any sequence of homotopy moves that tightens the interior tangle with strands α and β. Any $2\to0$ move involving both α and β removes one positive and one negative vertex; no other homotopy move changes the number of vertices in $\alpha \cap \beta$ or the signs of those vertices. Thus, tightening α and β leaves the sum $\sum_{x \in \alpha \cap \beta} \operatorname{sgn}(x)$ unchanged. Similarly, tightening the exterior tangle $\delta \cup \varepsilon$ leaves the sum $\sum_{y \in \delta \cap \varepsilon} \operatorname{sgn}(y)$ unchanged. But after tightening both tangles, either α and β are disjoint, or δ and ε are disjoint, or both, as γ is a single closed curve. Thus, at least one of the sums $\sum_{x \in \alpha \cap \beta} \operatorname{sgn}(x)$ and $\sum_{y \in \delta \cap \varepsilon} \operatorname{sgn}(y)$ is equal to zero. We conclude that mixed vertex pairs do not contribute to the defect.

The curve $\gamma \Cap \sigma$ obtained by tightening α and β has at most one interior vertex (and therefore no interior vertex pairs); the exterior vertices of $\gamma \Cap \sigma$ are precisely the exterior vertices of γ. Similarly, the curve $\gamma \Cup \sigma$ obtained by tightening both δ and ε has at most one exterior vertex; the interior vertices of $\gamma \Cup \sigma$ are precisely the interior vertices of γ. It follows that *defect*(γ) = *defect*$(\gamma \Cup \sigma)$ + *defect*$(\gamma \Cap \sigma)$.

[2] We recommend pronouncing $\Cap$ as "tightened inside" and $\Cup$ as "tightened outside"; note that the symbols $\Cap$ and $\Cup$ resemble the second letters of "inside" and "outside".

Finally, let γ' be the result of flipping the interior tangle. The curve $\gamma' \Cup \sigma$ is just a reflection of $\gamma \Cup \sigma$, which implies that $\mathit{defect}(\gamma' \Cup \sigma) = \mathit{defect}(\gamma \Cup \sigma)$, and straightforward case analysis implies $\gamma' \Cap \sigma = \gamma \Cap \sigma$. We conclude that $\mathit{defect}(\gamma') = \mathit{defect}(\gamma' \Cap \sigma) + \mathit{defect}(\gamma' \Cup \sigma) = \mathit{defect}(\gamma \Cap \sigma) + \mathit{defect}(\gamma \Cup \sigma) = \mathit{defect}(\gamma)$. □

Lemmas 5.3, 5.4, and 5.5 now immediately imply Theorem 5.1.

5.4 Back to planar electrical moves

Each planar electrical transformation in a plane graph G induces the same change in the medial graph $G^\times$ as a finite sequence of 1- and 2-strand tangle flips (hereafter simply called "tangle flips") followed by a single electrical move. For an arbitrary connected multicurve γ, let $\boldsymbol{\bar{X}(\gamma)}$ denote the minimum number of electrical moves in a mixed sequence of electrical moves and tangle flips that tightens γ. Similarly, let $\boldsymbol{\bar{H}(\gamma)}$ denote the minimum number of homotopy moves in a mixed sequence of homotopy moves and tangle flips that tightens γ. We emphasize that tangle flips are "free" and do not contribute to either $\bar{X}(\gamma)$ or $\bar{H}(\gamma)$.

Our lower bound on planar electrical moves follows our earlier lower bound proof for facial electrical moves almost verbatim; the only subtlety is that the embedding of the graph can effectively change at every step of the reduction. We repeat the arguments here to keep the presentation self-contained.

Lemma 5.6. $\bar{X}(\check{\gamma}) \le \bar{X}(\gamma)$ *for every connected proper smoothing* $\check{\gamma}$ *of every connected multicurve* γ *on the sphere.*

Proof: Let γ be a connected multicurve, and let $\check{\gamma}$ be a connected proper smoothing of γ. The proof proceeds by induction on $\bar{X}(\gamma)$. If $\bar{X}(\gamma) = 0$, then γ is already tight, so the lemma is vacuously true.

First, suppose $\check{\gamma}$ is obtained from γ by smoothing a single vertex x. Consider an optimal mixed sequence of tangle flips and electrical moves that tightens γ. This sequence starts with zero or more tangle flips, followed by a electrical move. Let γ' be the multicurve that results from the initial sequence of tangle flips; by definition, we have $\bar{X}(\gamma) = \bar{X}(\gamma')$. Moreover, applying the same sequence of tangle flips to $\check{\gamma}$ yields a connected multicurve $\check{\gamma}'$ such that $\bar{X}(\check{\gamma}) = \bar{X}(\check{\gamma}')$. Thus, we can assume without loss of generality that the first operation in the sequence is an electrical move.

Now let γ' be the result of this move; by definition, we have $\bar{X}(\gamma) = \bar{X}(\gamma') + 1$. As in the proof of Lemma 4.5, there are several subcases to consider, depending on whether the move from γ to γ' involves the smoothed vertex x, and if so, the specific type of move. In every subcase, by Lemma 3.1 we can apply at most one electrical move to $\check{\gamma}$ to obtain a (possibly trivial) smoothing $\check{\gamma}'$ of γ', and then apply the inductive hypothesis on γ' and $\check{\gamma}'$ to prove the statement. We omit the straightforward details.

Finally, if $\check{\gamma}$ is obtained from γ by smoothing more than one vertex, the lemma follows immediately by induction from the previous analysis. □

Lemma 5.7. *For every connected multicurve* γ*, there is an intermixed sequence of electrical moves and tangle flips that tightens* γ *that contains exactly* $\bar{X}(\gamma)$ *electrical moves, and does not contain* $0{\to}1$ *or* $1{\to}2$ *moves.*

Proof: Consider an optimal sequence of electrical moves and tangle flips that tightens γ, and let γ_i denote the result of the first i moves in this sequence. If any γ_i has more vertices than its predecessor γ_{i-1}, then γ_{i-1} is a connected proper smoothing of γ_i, and Lemma 5.6 implies a contradiction. □

Lemma 5.8. $\bar{X}(\gamma) \ge \bar{H}(\gamma)$ *for every closed curve* γ *on the sphere.*

Proof: Let γ be a closed curve on the sphere. The proof proceeds by induction on $\bar{X}(\gamma)$. If $\bar{X}(\gamma) = 0$, then γ is simple and thus $\bar{H}(\gamma) = 0$, so assume otherwise.

Consider an optimal sequence of electrical moves and tangle flips that tightens γ, and let γ_i be the curve obtained by applying a prefix of the sequence up to and including the first electrical move. The minimality of the sequence implies that $\bar{X}(\gamma) = \bar{X}(\gamma') + 1$. By Lemma 5.7, we can assume without loss of generality that the first electrical move in the sequence is neither 0→1 nor 1→2, and if this first electrical move is 1→0 or 3→3, the theorem immediately follows by induction.

The only remaining move to consider is 2→1. Let γ° denote the result of applying the same sequence of tangle flips to γ, but replacing the final 2→1 move with a 2→0 move, or equivalently, smoothing the vertex of γ' left by the final 2→1 move. We immediately have $\bar{H}(\gamma) \leq \bar{H}(\gamma^\circ) + 1$. Because γ° is a connected proper smoothing of γ', Lemma 5.6 implies $\bar{X}(\gamma^\circ) < \bar{X}(\gamma') = \bar{X}(\gamma) - 1$. Finally, the inductive hypothesis implies that $\bar{X}(\gamma^\circ) \geq \bar{H}(\gamma^\circ)$, which completes the proof. □

Lemma 5.9. *$\bar{H}(\gamma) \geq |defect(\gamma)|/2$ for every closed curve γ on the sphere.*

Proof: Each homotopy move decreases $|defect(\gamma)|$ by at most 2, and Lemmas 5.4 and 5.5 imply that tangle flips do not change $|defect(\gamma)|$ at all. Every simple curve has defect 0. □

Theorem 5.10. *Let G be an arbitrary planar graph, and let γ be any unicursal smoothing of $G^\times$ (defined with respect to any planar embedding of G). Reducing G to a single vertex requires at least $|defect(\gamma)|/2$ planar electrical transformations.*

Proof: The minimum number of planar electrical transformations required to reduce G is at least $\bar{X}(G^\times)$. Because γ is a single curve, it must be connected, so Lemma 5.6 implies that $\bar{X}(G^\times) \geq \bar{X}(\gamma)$. The theorem now follows immediately from Lemmas 5.8 and 5.9. □

Finally, Hayashi *et al.* [29] and Even-Zohar *et al.* [18] describe infinite families of planar closed curves with defect $\Omega(n^{3/2})$; see also [10, Section 2.2].

Corollary 5.11. *Reducing any n-vertex planar graph to a single vertex requires $\Omega(n^{3/2})$ planar electrical transformations in the worst case.*

6 Open problems

Our results suggest several open problems. Perhaps the most compelling, and the primary motivation for our work, is to find either a subquadratic upper bound or a quadratic lower bound on the number of (unrestricted) electrical transformations required to reduce any planar graph without terminals to a single vertex. Like Gitler [21], Feo and Provan [20], and Archdeacon *et al.* [3], we conjecture that $O(n^{3/2})$ *facial* electrical transformations suffice. However, proving the conjecture appears to be challenging.

Another direction is to prove a quadratic lower bound for graphs on surfaces with positive genus under *crossing-free* electrical transformations. To generalize Theorem 5.1 to surface-embedded graphs, we need an extension of Lemma 5.3 to navigate through all the possible embeddings. Using the theory of *large-edgewidth (LEW) embeddings*, a result by Thomassen [48, Theorem 6.1] shows that any embedding of a surface-embedded graph can be obtained from the LEW-embedding (if there's one) by a finite sequence of split reflections. From here it is not hard to construct a toroidal curve that admits an LEW-embedding and has quadratic defect. The main difficulty is that we don't have a similar electrical-homotopy inequality for arbitrary surfaces.

Finally, none of our lower bound techniques imply anything about non-planar electrical transformations or about electrical reduction of non-planar graphs. Indeed, the only lower bound known in the most general setting, for *any* family of electrically reducible graphs, is the trivial $\Omega(n)$. It seems unlikely that planar graphs can be reduced more quickly by using non-planar electrical transformations, but we can't prove anything. Any non-trivial lower bound for this problem would be interesting.

References

[1] Virgil W. Adkisson. Cyclicly connected continuous curves whose complementary domain boundaries are homeomorphic, preserving branch points. *C. R. Séances Soc. Sci. Lett. Varsovie III* 23:164–193, 1930.

[2] Francesca Aicardi. Tree-like curves. *Singularities and Bifurcations*, 1–31, 1994. Advances in Soviet Mathematics 21, Amer. Math. Soc.

[3] Dan Archdeacon, Charles J. Colbourn, Isidoro Gitler, and J. Scott Provan. Four-terminal reducibility and projective-planar wye-delta-wye-reducible graphs. *J. Graph Theory* 33(2):83–93, 2000.

[4] Vladimir I. Arnold. Plane curves, their invariants, perestroikas and classifications. *Singularities and Bifurcations*, 33–91, 1994. Adv. Soviet Math. 21, Amer. Math. Soc.

[5] Vladimir I. Arnold. *Topological Invariants of Plane Curves and Caustics*. University Lecture Series 5. Amer. Math. Soc., 1994.

[6] Jaizhen Cai. Counting embeddings of planar graphs using DFS trees. *SIAM J. Discrete Math.* 6(3):335–352, 1993.

[7] Gregory R. Chambers and Yevgeny Liokumovich. Converting homotopies to isotopies and dividing homotopies in half in an effective way. *Geometric and Functional Analysis* 24(4):1080–1100. Springer, 2014.

[8] Hsien-Chih Chang. *Tightening curves and graphs on surfaces*. Ph.D. dissertation, University of Illinois at Urbana-Champaign, 2018.

[9] Hsien-Chih Chang and Jeff Erickson. Electrical reduction, homotopy moves, and defect. Preprint, October 2015. arXiv:1510.00571.

[10] Hsien-Chih Chang and Jeff Erickson. Untangling planar curves. *Discrete & Computational Geometry* 58(4):889–920, 2017.

[11] Hsien-Chih Chang and Jeff Erickson. Unwinding annular curves and electrically reducing planar networks. Accepted to Computational Geometry: Young Researchers Forum, Proc. 33rd Int. Symp. Comput. Geom., 2017.

[12] Hsien-Chih Chang, Jeff Erickson, Arnaud de Mesmay, David Letscher, Saul Schleimer, Eric Sedgwick, Dylan Thurston, and Stephan Tillmann. Untangling curves on surfaces via local moves. *Proc. 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, 121–135, 2018.

[13] Yves Colin de Verdière, Isidoro Gitler, and Dirk Vertigan. Réseaux électriques planaires II. *Comment. Math. Helvetici* 71:144–167, 1996.

[14] Lino Demasi and Bojan Mohar. Four terminal planar Delta-Wye reducibility via rooted $K_{2,4}$ minors. *Proc. 26th Ann. ACM-SIAM Symp. Discrete Algorithms*, 1728–1742, 2015.

[15] Giuseppe Di Battista and Roberto Tamassia. Incremental planarity testing. *Proc. 30th Ann. IEEE Symp. Foundations Comput. Sci.*, 436–441, 1989.

[16] Giuseppe Di Battista and Roberto Tamassia. On-line planarity testing. *SIAM J. Comput.* 25(5):956–997, 1996.

[17] G. V. Epifanov. Reduction of a plane graph to an edge by a star-triangle transformation. *Dokl. Akad. Nauk SSSR* 166:19–22, 1966. In Russian. English translation in *Soviet Math. Dokl.* 7:13–17, 1966.

[18] Chaim Even-Zohar, Joel Hass, Nati Linial, and Tahl Nowik. Invariants of random knots and links. *Discrete & Computational Geometry* 56(2):274–314, 2016. arXiv:1411.3308.

[19] Thomas A. Feo. *I. A Lagrangian Relaxation Method for Testing The Infeasibility of Certain VLSI Routing Problems. II. Efficient Reduction of Planar Networks For Solving Certain Combinatorial Problems*. Ph.D. thesis, Univ. California Berkeley, 1985. ⟨http://search.proquest.com/docview/303364161⟩.

[20] Thomas A. Feo and J. Scott Provan. Delta-wye transformations and the efficient reduction of two-terminal planar graphs. *Oper. Res.* 41(3):572–582, 1993.

[21] Isidoro Gitler. *Delta-wye-delta Transformations: Algorithms and Applications*. Ph.D. thesis, Department of Combinatorics and Optimization, University of Waterloo, 1991.

[22] Isidoro Gitler and Feliú Sagols. On terminal delta-wye reducibility of planar graphs. *Networks* 57(2):174–186, 2011.

[23] Maurits de Graaf. *Graphs and curves on surfaces*. Ph.D. dissertation, Universiteit van Amsterdam, 1994.

[24] Maurits de Graaf and Alexander Schrijver. Characterizing homotopy of systems of curves on a compact surface by crossing numbers. *Linear Alg. Appl.* 226–228:519–528, 1995.

[25] Maurits de Graaf and Alexander Schrijver. Decomposition of graphs on surfaces. *J. Comb. Theory Ser. B* 70:157–165, 1997.

[26] Maurits de Graaf and Alexander Schrijver. Making curves minimally crossing by Reidemeister moves. *J. Comb. Theory Ser. B* 70(1):134–156, 1997.

[27] Joel Hass and Peter Scott. Intersections of curves on surfaces. *Israel J. Math.* 51:90–120, 1985.

[28] Joel Hass and Peter Scott. Shortening curves on surfaces. *Topology* 33(1):25–43, 1994.

[29] Chuichiro Hayashi, Miwa Hayashi, Minori Sawada, and Sayaka Yamada. Minimal unknotting sequences of Reidemeister moves containing unmatched RII moves. *J. Knot Theory Ramif.* 21(10):1250099 (13 pages), 2012. arXiv:1011.3963.

[30] Heinz Hopf. Über die Drehung der Tangenten und Sehnen ebener Kurven. *Compositio Math.* 2:50–62, 1935.

[31] Arthur Edwin Kennelly. Equivalence of triangles and three-pointed stars in conducting networks. *Electrical World and Engineer* 34(12):413–414, 1899.

[32] Saunders Mac Lane. A structural characterization of planar combinatorial graphs. *Duke Math. J.* 3(3):460–472, 1937.

[33] Hiroyuki Nakahara and Hiromitsu Takahashi. An algorithm for the solution of a linear system by Δ-Y transformations. *IEICE TRANSACTIONS on Fundamentals of Electronics, Communications and Computer Sciences* E79-A(7):1079–1088, 1996. Special Section on Multi-dimensional Mobile Information Network.

[34] Max Neumann-Coto. A characterization of shortest geodesics on surfaces. *Algebraic & Geometric Topology* 1:349–368, 2001.

[35] Steven D. Noble and Dominic J. A. Welsh. Knot graphs. *J. Graph Theory* 34(1):100–111, 2000.

[36] Michael Polyak. Invariants of curves and fronts via Gauss diagrams. *Topology* 37(5):989–1009, 1998.

[37] Neil Robertson and Paul D. Seymour. Graph minors. VII. Disjoint paths on a surface. *J. Comb. Theory Ser. B* 45(2):212–254, 1988.

[38] Neil Robertson and Richard Vitray. Representativity of surface embeddings. *Paths, Flows, and VLSI-Layout*, 293–328, 1990. Algorithms and Combinatorics 9, Springer-Verlag.

[39] Alexander Schrijver. Homotopy and crossing of systems of curves on a surface. *Linear Alg. Appl.* 114–115:157–167, 1989.

[40] Alexander Schrijver. Decomposition of graphs on surfaces and a homotopic circulation theorem. *J. Comb. Theory Ser. B* 51(2):161–210, 1991.

[41] Alexander Schrijver. Circuits in graphs embedded on the torus. *Discrete Math.* 106/107:415–433, 1992.

[42] Alexander Schrijver. On the uniqueness of kernels. *J. Comb. Theory Ser. B* 55:146–160, 1992.

[43] Matthias F. M. Stallmann. Using PQ-trees for planar embedding problems. Tech. Rep. NCSU-CSC TR-85-24, Dept. Comput. Sci., NC State Univ., December 1985. ⟨https://people.engr.ncsu.edu/mfms/Publications/1985-TR_NCSU_CSC-PQ_Trees.pdf⟩.

[44] Matthias F. M. Stallmann. On counting planar embeddings. *Discrete Math.* 122:385–392, 1993.

[45] Ernst Steinitz. Polyeder und Raumeinteilungen. *Enzyklopädie der mathematischen Wissenschaften mit Einschluss ihrer Anwendungen* III.AB(12):1–139, 1916.

[46] Ernst Steinitz and Hans Rademacher. *Vorlesungen über die Theorie der Polyeder: unter Einschluß der Elemente der Topologie*. Grundlehren der mathematischen Wissenschaften 41. Springer-Verlag, 1934. Reprinted 1976.

[47] Peter Guthrie Tait. On knots I. *Proc. Royal Soc. Edinburgh* 28(1):145–190, 1876–7.

[48] Carsten Thomassen. Embeddings of graphs with no short noncontractible cycles. *J. Comb. Theory Ser. B* 48(2):155–177, 1990.

[49] Klaus Truemper. On whitney's 2-isomorphism theorem for graphs. *Journal of Graph Theory* 4(1):43–49, 1980.

[50] Klaus Truemper. On the delta-wye reduction for planar graphs. *J. Graph Theory* 13(2):141–148, 1989.

[51] Klaus Truemper. *Matroid Decomposition*. Academic Press, 1992.

[52] Donald Wagner. Delta-wye reduction of almost-planar graphs. *Discrete Appl. Math.* 180:158–167, 2015.

[53] Hassler Whitney. Congruent graphs and the connectivity of graphs. *Amer. J. Math.* 54(1):150–168, 1932.

[54] Hassler Whitney. 2-isomorphic graphs. *American Journal of Mathematics* 55(1):245–254, 1933.

A Equivalence between electrical and homotopic tightness for primitive curves

A closed curve γ is ***primitive*** if γ is not homotopic to a proper multiple of some other closed curve. A multicurve is *primitive* if all its constituent curves are primitive. We show a five-way equivalence between electrical and homotopic tightness for primitive multicurves, which is implicit in the work by de Graaf and Schrijver [24, 25, 26, 39, 40, 41, 42].

Let γ be a multicurve on an orientable surface Σ such that each constituent curve of γ is primitive. Define the ***μ-function*** as

$$\boldsymbol{\mu(\gamma,\sigma)} := \min_{\substack{\sigma'\sim\sigma \\ \sigma'\pitchfork\gamma}} \operatorname{cr}(\gamma,\sigma'),$$

where $\operatorname{cr}(\gamma,\sigma')$ is the number of crossings between γ and σ', and the minimum ranges over every closed curve σ' homotopic to the given closed curve σ on Σ, intersecting γ transversely.[3] Denote $\boldsymbol{\mu_\gamma}$ as the single-variable function $\mu(\gamma,\cdot)$. The notion of μ-function is deeply related to the *representativity* or *facewidth* of a graph studied in topological graph theory [37, 38, 48].

The μ-function is a higher-genus analogue to the *depth* function defined in the annulus (see Section 4.1); in particular, both μ and *depth* are invariant under isotopy of γ and the electrical moves [38].

Lemma A.1 (Robertson and Vitray [38, Proposition 14.4]). *Electrical moves do not change μ_γ for any multicurve γ on surface Σ.*

Proof: For any face of γ intersected by some closed curve σ that could be deleted after an electrical move, exhaustive case analysis implies that there is another closed curve σ' that avoids that face. □

Multicurve γ satisfies ***simplicity conditions*** [40] if (1) any lifting of γ_i in the universal cover $\hat{\Sigma}$ does not self-intersect for any constituent curve γ_i of γ, and (2) any distinct liftings of γ_i and γ_j in $\hat{\Sigma}$ intersect each other at most once for any pair of (possibly identical) constituent curves γ_i and γ_j of γ. Multicurve γ is ***minimally crossing*** [40, 42] if each constituent curve of γ has minimum number of self-intersections in its homotopy class, and every pair of constituent curves has minimum intersections with each other, in their own homotopy classes. In notation, one has

$$\operatorname{cr}(\gamma_i) = \min_{\gamma_i'\sim\gamma_i} \operatorname{cr}(\gamma_i') \qquad \text{and} \qquad \operatorname{cr}(\gamma_i,\gamma_j) = \min_{\substack{\gamma_i'\sim\gamma_i \\ \gamma_j'\sim\gamma_j}} \operatorname{cr}(\gamma_i',\gamma_j')$$

for all constituent curves γ_i and γ_j of γ; $\operatorname{cr}(\gamma_i)$ denotes the number of self-intersections of curve γ_i. Multicurve γ is ***crossing-tight*** [40, 42] if $\mu_\gamma \neq \mu_{\check{\gamma}}$ for any proper smoothing $\check{\gamma}$ of γ.

Our proof of equivalence relies on machineries developed extensively in the sequence of work by de Graaf and Schrijver [24, 25, 26, 39, 40, 41, 42] who did all the weight-lifting. However the original work does not address the problem of relating electrical and homotopy moves.

Theorem A.2. *Let γ be a multicurve on an orientable surface whose constituent curves are all primitive. The following statements are equivalent: (1) Multicurve γ satisfies simplicity conditions, (2) γ is minimally crossing, (3) γ is crossing-tight, (4) γ is e-tight, and (5) γ is h-tight.*

[3] In Schrijver [42], the μ-function is defined with respect to the graph corresponding to γ through medial construction; the function defined here is denoted as μ' in his paper.

Proof: (1) $\Longleftrightarrow$ (2) $\Longleftrightarrow$ (3): Schrijver [40, Proposition 12] showed that γ satisfies simplicity conditions if and only if γ is minimally crossing and each constituent curve is primitive. Later in the same paper [40, Theorem 5] he also showed that γ is minimally crossing and each constituent curve is primitive if and only if γ is crossing-tight. An alternative proof using the monotonicity of homotopy process can be found in de Graaf's thesis [23].

(3) $\Rightarrow$ (4): In another paper Schrijver [42, Theorem 2] showed that two crossing-tight multicurves γ and γ' can be transformed into each other using only 3→3 moves if (and only if) $\mu_\gamma = \mu_{\check{\gamma}}$. This result implies that if multicurve γ is crossing-tight then γ is e-tight, as electrical moves preserves the μ-function by Lemma A.1.

(4) $\Rightarrow$ (5): Any e-tight multicurve must be h-tight by de Graaf and Schrijver [26] (see Lemma 3.5).

(5) $\Rightarrow$ (1): If γ is h-tight and primitive, then by Hass and Scott [27, Lemma 3.4] multicurve γ satisfies simplicity conditions. To elaborate, assume for contradiction that γ violates the simplicity conditions. As γ is h-tight one can push each constituent curve of γ close to its unique geodesic on the surface without even decreases the number of vertices, similar to the algorithm of de Graaf and Schrijver [26]. Therefore all the intersections between lifts of constituent curves of γ remains after the push. The primitiveness of the curve γ guarantees that each lift of any constituent curve does not self-intersect, and two different lifts of the same constituent curve intersects at most once on $\hat{\Sigma}$. Between the lifts of two distinct geodesics there is at most one intersection in the universal cover, and thus the same holds for the lifts of two distinct constituent curves of γ. This concludes the proof. $\square$

B Proving Lemma 3.1

Proof: We prove the statement by induction on the number of electrical moves in the sequence and the number of smoothed vertices. If $\check{\gamma} = \gamma$ then the statement trivially holds. Otherwise, we first consider the special case where $\check{\gamma}$ is obtained from γ by smoothing a single vertex x. Without loss of generality let γ' be the result of the first electrical move. There are two nontrivial cases to consider.

First, suppose the move from γ to γ' does not involve the smoothed vertex x. Then we can apply the same move to $\check{\gamma}$ to obtain a new multicurve $\check{\gamma}'$; the same multicurve can also be obtained from γ' by smoothing x.

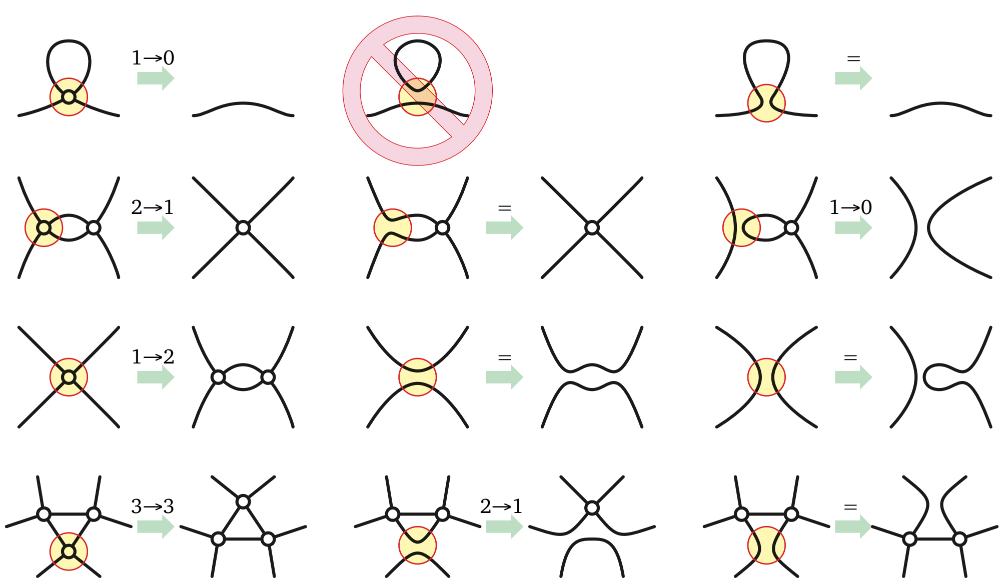

Figure B.1. Cases for the proof of the Lemma 3.1; the circled vertex is x.

Now suppose the first move does involve x. In this case, we can apply at most one electrical move to $\check{\gamma}$ to obtain a (possibly trivial) smoothing $\check{\gamma}'$ of γ'. There are eight subcases to consider, shown in Figure B.1. One subcase for the 1→0 move is impossible, because $\check{\gamma}$ is connected. In the remaining 1→0 subcase and one 2→1 subcase, the curves $\check{\gamma}$, $\check{\gamma}'$, and γ' are all isomorphic. In all remaining subcases, $\check{\gamma}'$ is a connected proper smoothing of γ'.

Finally, we consider the more general case where $\check{\gamma}$ is obtained from γ by smoothing more than one vertex. Let $\tilde{\gamma}$ be any intermediate curve, obtained from γ by smoothing just one of the vertices that were smoothed to obtain $\check{\gamma}$. As $\check{\gamma}$ is a connected smoothing of $\tilde{\gamma}$, the curve $\tilde{\gamma}$ itself must be connected too. Our earlier argument implies that there is a sequence of electrical moves that changes $\tilde{\gamma}$ to a smoothing $\tilde{\gamma}'$ of γ'. The inductive hypothesis implies that there is a sequence of electrical moves that changes $\check{\gamma}$ to a smoothing $\check{\gamma}'$ of $\tilde{\gamma}'$, which is itself a smoothing of γ'. This completes the proof. □