Untangling Planar Curves*

Hsien-Chih Chang[1] and Jeff Erickson[1]

1 **Department of Computer Science**
University of Illinois, Urbana-Champaign
{hchang17, jeffe}@illinois.edu

Abstract

Any generic closed curve in the plane can be transformed into a simple closed curve by a finite sequence of local transformations called homotopy moves. We prove that simplifying a planar closed curve with n self-crossings requires $\Theta(n^{3/2})$ homotopy moves in the worst case. Our algorithm improves the best previous upper bound $O(n^2)$, which is already implicit in the classical work of Steinitz; the matching lower bound follows from the construction of closed curves with large *defect*, a topological invariant of generic closed curves introduced by Aicardi and Arnold. This lower bound also implies that $\Omega(n^{3/2})$ degree-1 reductions, series-parallel reductions, and ΔY transformations are required to reduce any planar graph with treewidth $\Omega(\sqrt{n})$ to a single edge, matching known upper bounds for rectangular and cylindrical grid graphs. Finally, we prove that $\Omega(n^2)$ homotopy moves are required in the worst case to transform one non-contractible closed curve on the torus to another; this lower bound is tight if the curve is homotopic to a simple closed curve.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases computational topology, homotopy, planar graphs, ΔY transformations, defect, Reidemeister moves, tangles

Digital Object Identifier 10.4230/LIPIcs.SoCG.2016.29

1 Introduction

Any regular closed curve in the plane can be transformed into a simple closed curve by a finite sequence of the following local operations:

- 1→0: Remove an empty loop.
- 2→0: Separate two subpaths that bound an empty bigon.
- 3→3: Flip an empty triangle by moving one subpath over the opposite intersection point.

See Figure 1.1. Each of these operations can be performed by continuously deforming the curve within a small neighborhood of one face; consequently, we call these operations and their inverses *homotopy moves*. Our notation is nonstandard but mnemonic; the numbers before and after each arrow indicate the number of local vertices before and after the move. Homotopy moves are "shadows" of the classical Reidemeister moves used to manipulate knot and link diagrams [4, 36].

We prove that $\Theta(n^{3/2})$ homotopy moves are required in the worst case to simplify a closed curve in the plane with n self-crossings. Before describing our results in more detail, we review several previous results.

* Work on this paper was partially supported by NSF grant CCF-1408763. See http://jeffe.cs.illinois.edu/pubs/tangle.pdf for the most recent version of this paper.

32nd International Symposium on Computational Geometry (SoCG 2016).
Editors: Sándor Fekete and Anna Lubiw; Article No. 29; pp. 29:1–29:15
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

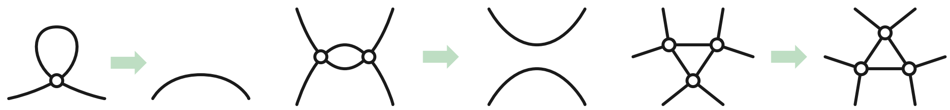

Figure 1.1 Homotopy moves 1→0, 2→0, and 3→3.

1.1 Past Results

An algorithm to simplify any planar closed curve using at most $O(n^2)$ homotopy moves is implicit in Steinitz's proof that every 3-connected planar graph is the 1-skeleton of a convex polyhedron [41, 42]. Specifically, Steinitz proved that any non-simple closed curve (in fact, any 4-regular plane graph) with no empty loops contains a *lens* ("Spindel"): a disk bounded by a pair of simple subpaths that cross exactly twice, where the endpoints of the subpaths lie outside the disk. Steinitz then proves that any *minimal* lens ("irreduzible Spindel") can be transformed into an empty bigon with a sequence of 3→3 moves, each removing one triangular face from the lens, as shown in Figure 1.2. Once the lens is empty, it can be deleted with a single 2→0 move. See Grünbaum [23], Hass and Scott [25], Colin de Verdière *et al.* [12], or Nowik [33] for more modern treatments of Steinitz's technique. The $O(n^2)$ upper bound also follows from algorithms for *regular* homotopy, which forbids 0↔1 moves, by Francis [19], Vegter [45] (for polygonal curves), and Nowik [33].

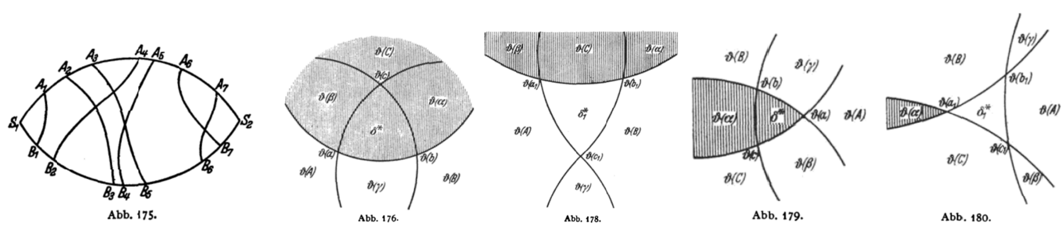

Figure 1.2 A minimal lens, and 3→3 moves removing triangles from the side or the end of a (shaded) minimal lens. All figures are from Steinitz and Rademacher [42].

The $O(n^2)$ upper bound can also be derived from an algorithm of Feo and Provan [17] for reducing a planar graph to a single edge by *electrical transformations*: degree-1 reductions, series-parallel reductions, and ΔY-transformations. (We consider electrical transformations in more detail in Section 3.) Any curve divides the plane into regions, called its *faces*. The *depth* of a face is its distance to the outer face in the dual graph of the curve. Call a homotopy move *positive* if it decreases the sum of the face depths; in particular, every 1→0 and 2→0 move is positive. Feo and Provan prove that every non-simple curve in the plane admits a positive homotopy move [17, Theorem 1]. Thus, the sum of the face depths is an upper bound on the minimum number of moves required to simplify the curve. Euler's formula implies that every curve with n crossings has $O(n)$ faces, and each of these faces has depth $O(n)$.

Gitler [21] conjectured that a variant of Feo and Provan's algorithm that always makes the *deepest* positive move requires only $O(n^{3/2})$ moves. Song [40] observed that if Feo and Provan's algorithm always chooses the *shallowest* positive move, it can be forced to make $\Omega(n^2)$ moves even when the input curve can be simplified using only $O(n)$ moves.

Tight bounds are known for two special cases that forbid certain types of homotopy moves. First, Nowik proved a tight $\Omega(n^2)$ lower bound for regular homotopy [33]. Second, Khovanov [30] defined two curves to be *doodle equivalent* if one can be transformed into the other using 1↔0 and 2↔0 moves. Khovanov [30] and Ito and Takimura [28] independently proved that any planar curve can be transformed into its unique equivalent doodle with the smallest number of vertices, using only 1→0 and 2→0 moves. Thus, two doodle equivalent curves are connected by a sequence of $O(n)$ moves, which is obviously tight.

1.2 New Results

In Section 2, we derive an $\Omega(n^{3/2})$ lower bound, using a numerical curve invariant called *defect*, introduced by Arnold [7, 8] and Aicardi [1]. Any homotopy move changes the defect of a closed curve by at most 2. Thus, the lower bound follows from constructions of Hayashi *et al.* [26, 27] and Even-Zohar *et al.* [16] of closed curves with defect $\Omega(n^{3/2})$. We simplify and generalize their results by computing the defect of the standard planar projection of any $p \times q$ torus knot where either $p \bmod q = 1$ or $q \bmod p = 1$. Our calculations imply that for any integer p, reducing the standard projection of the $p \times (p+1)$ torus knot requires at least $\binom{p+1}{3} \geq n^{3/2}/6 - O(n)$ homotopy moves.

In Section 3, we sketch a proof, based on arguments of Truemper [43] and Noble and Welsh [32], that reducing a planar graph G using electrical transformations requires at least as many steps as reducing the medial graph of G to a simple closed curve using homotopy moves. The homotopy lower bound from Section 2 then implies that reducing any n-vertex planar graph with treewidth $\Omega(\sqrt{n})$ requires $\Omega(n^{3/2})$ electrical transformations. This lower bound matches known upper bounds for rectangular and cylindrical grid graphs. Due to space limitations, we omit most technical details from this section; we refer the interested reader to our preprint [9].

We develop a new algorithm to simplify any closed curve in $O(n^{3/2})$ homotopy moves in Section 4. First we describe an algorithm that uses $O(D)$ moves, where D is the sum of the face depths of the input curve. At a high level, our algorithm can be viewed as a variant of Steinitz's algorithm that empties and removes *loops* instead of lenses. We then extend our algorithm to *tangles*: collections of boundary-to-boundary paths in a closed disk. Our algorithm simplifies a tangle as much as possible in $O(D + ns)$ moves, where D is the sum of the depths of the tangle's faces, s is the number of paths, and n is the number of intersection points. Finally, we prove that for any curve with maximum face depth $\Omega(\sqrt{n})$, we can find a simple closed curve whose interior tangle has $s = O(\sqrt{n})$ strands, maximum face depth $O(\sqrt{n})$, and at least s^2 interior vertices. Simplifying this tangle and then recursively simplifying the resulting curve requires a total of $O(n^{3/2})$ moves.

Finally, in Section 5, we consider the natural generalization of the homotopy problem to curves on higher-genus surfaces. We prove that $\Omega(n^2)$ homotopy moves are required in the worst case to transform one non-contractible closed curve on the torus to another. Results of Hass and Scott [24] imply that this lower bound is tight if the curve is homotopic to a simple closed curve.

1.3 Definitions

A *closed curve* in a surface M is a continuous map $\gamma \colon S^1 \to M$; in this paper, we consider only *regular* closed curves, which are injective except at a finite number of self-intersections, each of which is a transverse double point. A closed curve is *simple* if it is injective. For most of the paper, we consider only closed curves in the plane; we consider more general surfaces in Section 5.

The image of any non-simple closed curve has a natural structure as a 4-regular plane graph. Thus, we refer to the self-intersection points of a curve as its *vertices*, the maximal subpaths between vertices as *edges*, and the components of the complement of the curve as its *faces*. Two curves γ and γ' are isomorphic if their images are isomorphic as planar maps; we will not distinguish between isomorphic curves.

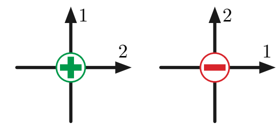

We adopt a standard sign convention for vertices first used by Gauss [20]. Choose an arbitrary *basepoint* $\gamma(0)$. We call a vertex *positive* if the first traversal through the vertex crosses the second traversal from right to left, and *negative* otherwise. We define $\operatorname{sgn}(x) = +1$ for every positive vertex x and $\operatorname{sgn}(x) = -1$ for every negative vertex x.

A *homotopy* between two curves γ and γ' is a continuous function $H\colon S^1 \times [0,1] \to M$ such that $H(\cdot, 0) = \gamma$ and $H(\cdot, 1) = \gamma'$. Each homotopy move can be executed by a homotopy. Conversely, Alexander's simplicial approximation theorem [3], together with combinatorial arguments of Alexander and Briggs [4] and Reidemeister [36], imply that any generic homotopy between two closed curves can be decomposed into a finite sequence of homotopy moves. Two curves are *homotopic*, or in the same *homotopy class*, if there is a homotopy from one to the other. All closed curves in the plane are homotopic.

2 Lower Bounds

2.1 Defect

To prove our main lower bound, we consider a numerical invariant of closed curves in the plane introduced by Arnold [7, 8] and Aicardi [1] called *defect*. Polyak [35] proved that defect can be computed—or for our purposes, defined—as follows:

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

Here the sum is taken over all *interleaved* pairs of vertices of γ: two vertices $x \neq y$ are interleaved, denoted $x \between y$, if they alternate in cyclic order—x, y, x, y—along γ. Even though the signs of individual vertices depend on the basepoint and orientation of the curve, the defect of a curve is independent of those choices. Trivially, every simple closed curve has defect zero. Straightforward case analysis [35] implies that any single homotopy move changes the defect of a curve by at most 2; the various cases are illustrated in Figure 2.1.

- A 1→0 move leaves the defect unchanged.

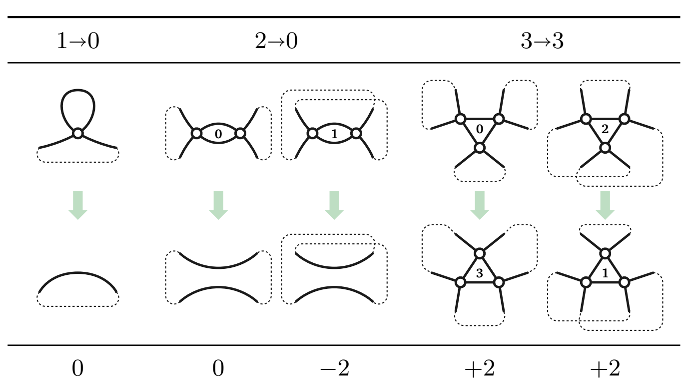

Figure 2.1 Changes to defect incurred by homotopy moves. Numbers in each figure indicate how many pairs of vertices are interleaved; dashed lines indicate how the rest of the curve connects.

- A 2→0 move decreases the defect by 2 if the two disappearing vertices are interleaved, and leaves the defect unchanged otherwise.
- A 3→3 move increases the defect by 2 if the three vertices before the move contain an even number of interleaved pairs, and decreases the defect by 2 otherwise.

▶ **Lemma 2.1.** *Let γ be an arbitrary closed curve in the plane. Simplifying γ requires at least $|defect(\gamma)|/2$ homotopy moves.*

2.2 Flat Torus Knots

For any relatively prime positive integers p and q, let $T(p,q)$ denote the curve with the following parametrization, where θ runs from 0 to 2π:

$$T(p,q)(\theta) := \big((\cos(q\theta)+2)\cos(p\theta),\ (\cos(q\theta)+2)\sin(p\theta)\big).$$

The curve $T(p,q)$ winds around the origin p times, oscillates q times between two concentric circles, and crosses itself exactly $(p-1)q$ times. We call these curves *flat torus knots*.

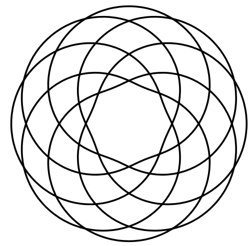

Figure 2.2 The flat torus knot $T(7,8)$.

Hayashi *et al.* [27, Proposition 3.1] proved that for any integer q, the flat torus knot $T(q+1,q)$ has defect $-2\binom{q}{3}$. Even-Zohar *et al.* [16] used a star-polygon representation of the curve $T(p,2p+1)$ as the basis for a universal model of random knots; in our notation, they proved that $defect(T(p,2p+1)) = 4\binom{p+1}{3}$ for any integer p. In this section we simplify and generalize both of these results to all $T(p,q)$ where either $q \bmod p = 1$ or $p \bmod q = 1$.

▶ **Lemma 2.2.** *$defect(T(p,ap+1)) = 2a\binom{p+1}{3}$ for all integers $a \geq 0$ and $p \geq 1$.*

Proof. For purposes of illustration, we cut any torus knot $T(p,q)$ open into a "flat braid" consisting of p x-monotone paths, which we call *strands*, between two fixed diagonal lines. All strands are directed from left to right.

The curve $T(p,1)$ can be reduced using only 1→0 moves, so its defect is zero. For any integer $a \geq 0$, we can reduce $T(p,ap+1)$ to $T(p,(a-1)p+1)$ by straightening the leftmost block of $p(p-1)$ crossings in the flat braid representation, one strand at a time. Within this block, each pair of strands in the flat braid intersect twice. Straightening the bottom strand of this block requires the following $\binom{p}{2}$ moves, as shown in Figure 2.3.

- $\binom{p-1}{2}$ 3→3 moves pull the bottom strand downward over one intersection point of every other pair of strands. Just before each 3→3 move, exactly one of the three pairs of the three relevant vertices is interleaved, so each move decreases the defect by 2.
- $(p-1)$ 2→0 moves eliminate a pair of intersection points between the bottom strand and every other strand. Each of these moves also decreases the defect by 2.

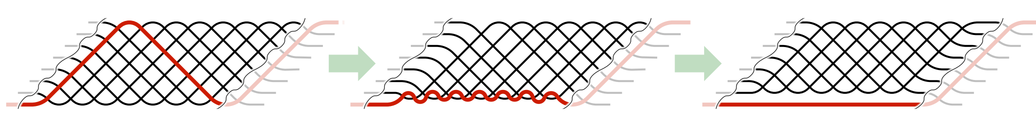

Figure 2.3 Straightening one strand in a block of $T(8,8a+1)$.

Altogether, straightening one strand decreases the defect by $2\binom{p}{2}$. Proceeding similarly with the other strands, we conclude that $defect(T(p,ap+1)) = defect(T(p,(a-1)p+1)) + 2\binom{p+1}{3}$. The lemma follows immediately by induction. ◀

A similar argument [9, Lemma 4.3] gives us the exact value of $defect(T(p,q))$ when $p \bmod q = 1$.

▶ **Lemma 2.3.** *$defect(T(aq+1,q)) = -2a\binom{q}{3}$ for all integers $a \geq 0$ and $q \geq 1$.*

Either of the previous lemmas imply the following lower bound, which is also implicit in Hayashi *et al.* [27].

▶ **Theorem 2.4.** *For every positive integer n, there is a closed curve with n crossings that requires at least $n^{3/2}/6 - O(n)$ homotopy moves to reduce to a simple closed curve.*

3 Electrical Transformations

Now we consider a related set of local operations on plane graphs, called *electrical transformations*, consisting of six operations in three dual pairs, as shown in Figure 3.1.

- *degree-1 reduction*: Contract the edge incident to a vertex of degree 1, or delete the edge incident to a face of degree 1
- *series-parallel reduction*: Contract either edge incident to a vertex of degree 2, or delete either edge incident to a face of degree 2
- ΔY *transformation*: Delete a vertex of degree 3 and connect its neighbors with three new edges, or delete the edges bounding a face of degree 3 and join the vertices of that face to a new vertex.

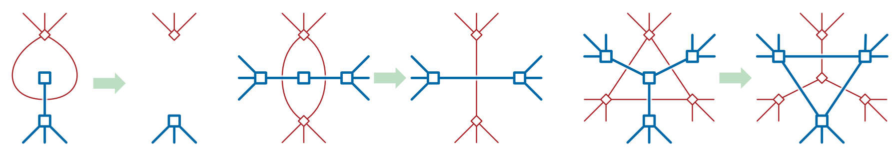

Figure 3.1 Electrical transformations in a plane graph G and its dual graph G^*.

Electrical transformations have been used since the end of the 19th century [29, 39] to analyze resistor networks and other electrical circuits, but have since been applied to a number of other combinatorial problems on planar graphs, including shortest paths and maximum flows [2]; multicommodity flows [18]; and counting spanning trees, perfect matchings, and cuts [11]. We refer to our earlier preprint [9, Section 1.1] for a more detailed history and an expanded list of applications.

Epifanov [15] was the first to prove that any planar graph with two special vertices called *terminals* can be reduced to a single edge between the terminals by a finite number of electrical transformations. Simpler algorithmic proofs were later given by Truemper [43], Feo and Provan [17], and others. In particular, Truemper's algorithm reduces the $p \times p$ grid using $O(p^3)$ moves, which is the best bound known for this special case. Since every n-vertex planar graph is a minor of a $\Theta(n) \times \Theta(n)$ grid [44], Truemper's algorithm implies an $O(n^3)$ upper bound for arbitrary planar graphs; see Lemma 3.1. Feo and Provan's algorithm uses $O(n^2)$ electrical transformations, which is the best general upper bound known.

However, for the simpler problem of reducing a planar graph without terminals to a single *vertex*, an $O(n^2)$-move algorithm already follows from the lens-reduction argument of Steinitz described in the introduction [41, 42]. Steinitz reduced local transformations of planar *graphs* to local transformations of planar *curves* by defining *medial graphs* ("Θ-Prozess").

The *medial graph* of a plane graph G, which we denote $G^\times$, is another plane graph whose vertices correspond to the edges of G and whose edges correspond to incidences (with multiplicity) between vertices of G and faces of G. Two vertices of $G^\times$ are connected by an edge if and only 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, and every 4-regular plane graph is a medial graph. A 4-regular plane graph is *unicursal* if it is the image of a single closed curve.

Electrical transformations in any plane graph G correspond to local transformations in the medial graph $G^\times$ that are almost identical to homotopy moves. Each degree-1 reduction in G corresponds to a 1→0 homotopy move in $G^\times$, and each ΔY transformation in G corresponds to a 3→3 homotopy move in $G^\times$. A series-parallel reduction in G contracts an empty bigon in $G^\times$ to a single vertex. Extending our earlier notation, we call this transformation a 2→1 move. We collectively refer to these transformations and their inverses as *medial electrical moves*; see Figure 3.2.

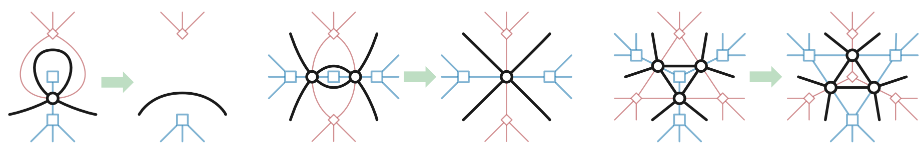

Figure 3.2 Electrical transformations and medial electrical moves 1→0, 2→1, and 3→3.

Here we sketch a proof that $\Omega(n^{3/2})$ electrical transformations are required in the worst case to reduce an n-vertex planar graph to a single vertex. Our proof builds on two key lemmas. The first lemma follows from close reading of the proofs by Truemper [43, Lemma 4] and Gitler [21, Lemma 2.3.3] that if a graph G can be can be reduced to a single vertex by electrical transformations, then so can every minor of G. The second lemma is implicit in the work of Noble and Welsh [32]; informally, we can replace 2→1 medial electrical moves with 2→0 homotopy moves. For self-contained proofs of these lemmas, which argue directly in terms of medial electrical moves, we refer to our earlier preprint [9, Section 3].

▶ **Lemma 3.1.** *Let G be any plane graph. Reducing any proper minor of G to a single vertex requires strictly fewer electrical transformations than reducing G to a single vertex.*

▶ **Lemma 3.2.** *Let G be any plane graph whose medial graph $G^\times$ is unicursal. The minimum number of homotopy moves required to reduce $G^\times$ to a simple closed curve is no greater than the minimum number of electrical transformations required to reduce G to a single vertex.*

▶ **Theorem 3.3.** *For every positive integer t, every planar graph with treewidth t requires $\Omega(t^3)$ electrical transformations to reduce to a single vertex.*

Proof. Every planar graph with treewidth t contains an $\Omega(t) \times \Omega(t)$ grid minor [38], which in turn contains an $k \times (2k+1)$ cylindrical grid minor for some integer $k = \Omega(t)$. The medial graph of the $k \times (2k+1)$ cylindrical grid is the flat torus knot $T(2k, 2k+1)$. The theorem now follows from Lemmas 2.1, 2.2, 3.1, and 3.2. ◀

In particular, reducing any planar graph with n vertices and treewidth $\Theta(\sqrt{n})$ requires $\Omega(n^{3/2})$ electrical transformations. It follows that Truemper's $O(p^3)$ upper bound for reducing the $p \times p$ square grid [43] is tight. Similar arguments using Lemmas 2.2 and 2.3 imply that Nakahara and Takahashi's $O(\min\{pq^2, p^2q\})$ upper bound for reducing the $p \times q$ cylindrical grid [31] is also tight.

Like Gitler [21], Feo and Provan [17], and Archdeacon *et al.* [6], we conjecture that any planar graph with n vertices can be reduced using only $O(n^{3/2})$ electrical transformations, but so far we have only been able to prove a matching upper bound for homotopy moves.

4 Upper Bound

For any point p, let $depth(p, \gamma)$ denote the minimum number of times a path from p to infinity crosses γ. Any two points in the same face of γ have the same depth, so each face f has a well-defined depth, which is its distance to the outer face in the dual graph of γ; see Figure 4.1. The depth of the curve, denoted $depth(\gamma)$, is the maximum depth of the faces of γ; and the *potential* $D(\gamma)$ is the sum of the depths of the faces. Euler's formula implies that any 4-regular planar graph with n vertices has exactly $n+2$ faces; thus, for any curve γ with n vertices, we have $n+1 \le D(\gamma) \le (n+1) \cdot depth(\gamma)$.

4.1 Contracting Simple Loops

▶ **Lemma 4.1.** *Every closed curve γ in the plane can be simplified using at most $3D(\gamma) - 3$ homotopy moves.*

Proof. The lemma is trivial if γ is already simple, so assume otherwise. Let $x \coloneqq \gamma(\theta) = \gamma(\theta')$ be the first vertex to be visited twice by γ after the (arbitrarily chosen) basepoint $\gamma(0)$. Let ℓ denote the subcurve of γ from $\gamma(\theta)$ to $\gamma(\theta')$; our choice of x implies that ℓ is a simple loop. Let m and s denote the number of vertices and maximal subpaths of γ in the interior of ℓ respectively. Finally, let γ' denote the closed curve obtained from γ by removing ℓ. The first stage of our algorithm transforms γ into γ' by contracting the loop ℓ via homotopy moves.

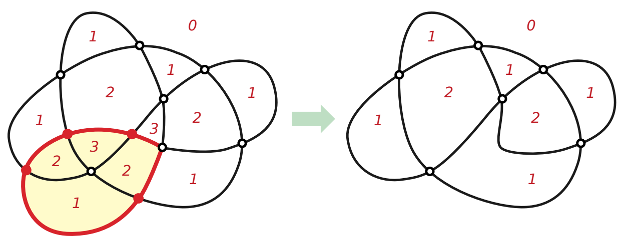

■ **Figure 4.1** Transforming γ into γ' by contracting a simple loop. Numbers are face depths.

We remove the vertices and edges from the interior of ℓ one at a time as follows. If we can perform a 2→0 move to remove one edge of γ from the interior of ℓ and decrease s, we do so. Otherwise, either ℓ is empty, or some vertex of γ lies inside ℓ. Let y be a vertex of γ that lies inside ℓ and has at least one neighbor z that lies on ℓ; we move y outside ℓ with a 0→2 move (which increases s) followed by a 3→3 move, as shown on the right of Figure 4.2.

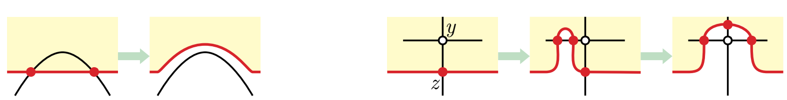

■ **Figure 4.2** Moving a loop over an interior empty bigon or an interior vertex.

Once ℓ is an empty loop, we remove it with a single 1→0 move. Altogether, our algorithm transforms γ into γ' using at most $3m + s + 1$ homotopy moves. Let M denote the actual number of homotopy moves used.

Euler's formula implies that ℓ contains exactly $m + s/2 + 1$ faces of γ. The Jordan Curve theorem implies that $depth(p, \gamma') \le depth(p, \gamma) - 1$ for any point p inside ℓ, and trivially $depth(p, \gamma') \le depth(p, \gamma)$ for any point p outside ℓ. It follows that $D(\gamma') \le D(\gamma) - (m + s/2 + 1) \le D(\gamma) - M/3$, and therefore $M \le 3D(\gamma) - 3D(\gamma')$. The induction hypothesis implies that we can recursively simplify γ' using at most $3D(\gamma') - 3$ moves. The lemma now follows immediately. ◀

Our upper bound is a factor of 3 larger than Feo and Provan's [17]; however our algorithm has the benefit that it extends to *tangles*, as described in the next subsection.

4.2 Tangles

A *tangle* is a collection of boundary-to-boundary paths $\gamma_1, \gamma_2, \ldots, \gamma_s$ in a closed topological disk Σ, which (self-)intersect only pairwise, transversely, and away from the boundary of Σ. (In knot theory, a tangle usually refers to the intersection of a knot or link with a closed 3-dimensional ball [10, 13]; our object is more properly called a *flat tangle*, as it is the image of a tangle under an appropriate projection. Our tangles are unrelated to the obstructions to small branchwidth studied by Robertson and Seymour [37].)

We call each individual path γ_i a *strand* of the tangle. The *boundary* of a tangle is the boundary of the disk Σ that contains it; we usually denote the boundary by σ. By the Jordan-Schönflies theorem, we can assume without loss of generality that σ is actually a circle. We can obtain a tangle from any closed curve γ by considering its restriction to any closed disk whose boundary σ intersects γ transversely away from its vertices; we call this restriction the *interior tangle* of σ.

The strands and boundary of any tangle define a plane graph T whose boundary vertices each have degree 3 and whose interior vertices each have degree 4. Depths and potential are defined exactly as for closed curves: The depth of any face f of T is its distance to the outer face in the dual graph T^*; the depth of the tangle is its maximum face depth; and the potential $D(T)$ of the tangle is the sum of its face depths.

A tangle is *tight* if every pair of strands intersects at most once and *loose* otherwise. Every loose tangle contains either an empty loop or a (not necessarily empty) lens. Thus, any tangle with n vertices can be transformed into a tight tangle—or less formally, *tightened*—in $O(n^2)$ homotopy moves using Steinitz's algorithm. On the other hand, there are infinite classes of loose tangles for which no homotopy move decreases the potential, so we cannot directly apply Feo and Provan's algorithm to this setting.

▶ **Lemma 4.2.** *Every tangle T with n vertices and s strands can be tightened using at most $3D(T) + 3ns$ homotopy moves.*

Proof. Our algorithm consists of two stages: first we simplify the individual strands using at most $3D(T)$ homotopy moves, and then we remove excess intersections between every pair of strands using at most $3ns$ homotopy moves.

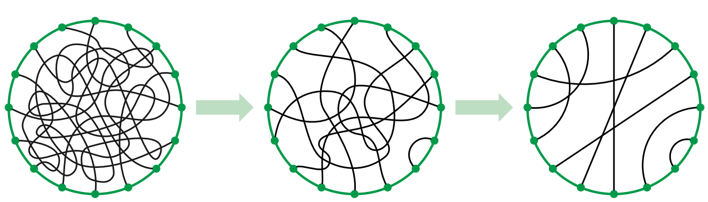

Figure 4.3 Tightening a tangle in two phases: First simplifying the individual strands, then removing excess crossings between pairs of strands.

First, as long as any strand in T is non-simple, we identify a simple loop ℓ in that strand and remove it as described in the proof of Lemma 4.1. Let T' be the remaining strand after all such loops are removed. The analysis in the proof of Lemma 4.1 implies that transforming T into T' requires at most $3D(T) - 3D(T') \le 3D(T)$ homotopy moves.

Now fix an arbitrary reference point on the boundary circle σ that is not an endpoint of a strand. For each index i, let σ_i be the arc of σ between the endpoints of γ_i that does not contain the reference point. A strand γ_i is *extremal* if the corresponding arc σ_i does not contain any other arc σ_j.

Choose an arbitrary extremal strand γ_i. Let m_i denote the number of tangle vertices in the interior of the disk bounded by γ_i and the boundary arc σ_i; let s_i denote the number of intersections between γ_i and other strands. Finally, let γ_i' be a path inside the disk Σ defining tangle T, with the same endpoints as γ_i, that intersects each other strand in T at most once, such that the disk bounded by σ_i and γ_i' has no tangle vertices inside its interior.

We can deform γ_i into γ_i' using essentially the algorithm from Lemma 4.1. If the disk bounded by γ_i and σ_i contains an empty bigon, remove it with a 2→0 move. If the disk has an interior vertex with a neighbor on γ_i, remove it using at most two homotopy moves (and possibly increasing s_i). Altogether, this deformation requires at most $3m_i + s_i \le 3n$ homotopy moves.

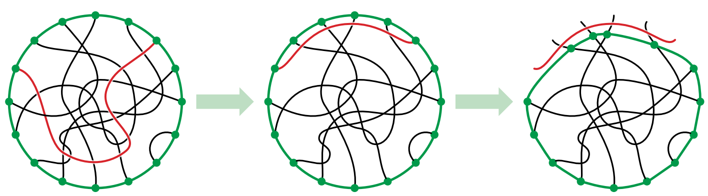

Figure 4.4 Moving one strand out of the way and shrinking the tangle boundary.

After deforming γ_i to γ_i', we shrink the boundary of the tangle slightly to exclude γ_i', without creating or removing any additional endpoints on the boundary or vertices in the tangle. We emphasize that shrinking the boundary does not modify the strands and therefore does not require any homotopy moves. The resulting smaller tangle has exactly $s-1$ strands, each of which is simple. Thus, the induction hypothesis implies that we can recursively tighten this smaller tangle using at most $3n(s-1)$ homotopy moves. The base case of this recursion is a tangle with no strands. ◀

4.3 Main Algorithm

We call a simple closed curve σ *useful* for γ if it intersects γ transversely away from its vertices, and the interior tangle of σ has at least s^2 vertices, where $s = |\sigma \cap \gamma|/2$ is the number of strands. Our main algorithm repeatedly finds a useful closed curve and tightens its interior tangle; if there are no useful closed curves, then we fall back to the loop-contraction algorithm of Lemma 4.1.

▶ **Lemma 4.3.** *Let γ be an arbitrary closed curve in the plane with n vertices. Either γ has a useful simple closed curve whose interior tangle has depth $O(\sqrt{n})$, or γ itself has depth $O(\sqrt{n})$.*

Proof. To simplify notation, let $d \coloneqq \mathit{depth}(\gamma)$. For each integer j between 1 and d, let R_j be the set of points p with $\mathit{depth}(p, \gamma) \ge d + 1 - j$, and let $\tilde{R}_j$ denote a small open neighborhood of the closure of $R_j \cup \tilde{R}_{j-1}$, where $\tilde{R}_0$ is the empty set. Each region $\tilde{R}_j$ is the disjoint union of closed disks, whose boundary cycles intersect γ transversely away from its vertices, if at all. In particular, $\tilde{R}_d$ is a disk containing the entire curve γ.

Fix a point z such that $depth(z, \gamma) = d$. For each integer j, let Σ_j be the unique component of $\tilde{R}_j$ that contains z, and let σ_j be the boundary of Σ_j. Then $\sigma_1, \sigma_2, \ldots, \sigma_d$ are disjoint, nested, simple closed curves; see Figure 4.5. Let n_j be the number of vertices and let $s_j \coloneqq |\gamma \cap \sigma_j|/2$ be the number of strands of the interior tangle of σ_j. For notational convenience, we define $\Sigma_0 \coloneqq \varnothing$ and thus $n_0 = s_0 = 0$.

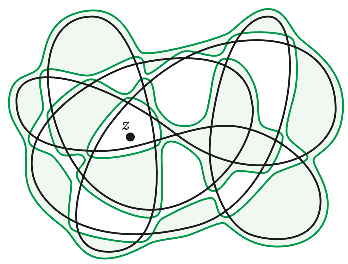

Figure 4.5 Nested depth cycles around a point of maximum depth.

By construction, for each j, the interior tangle of σ_j has depth $j + 1$. Thus, to prove the lemma, it suffices to show that either $depth(\gamma) = O(\sqrt{n})$ or at least one curve σ_j with $j = O(\sqrt{n})$ is useful.

Fix an index j. Each edge of γ crosses σ_j at most twice. Any edge of γ that crosses σ_j has at least one endpoint in the annulus $\Sigma_j \setminus \Sigma_{j-1}$, and any edge that crosses σ_j twice has both endpoints in $\Sigma_j \setminus \Sigma_{j-1}$. Conversely, each vertex in Σ_j is incident to at most two edges that cross σ_j and no edges that cross σ_{j+1}. It follows that $|\sigma_j \cap \gamma| \leq 2(n_j - n_{j-1})$, and therefore $n_j \geq n_{j-1} + s_j$. Thus, by induction, we have

$$n_j \geq \sum_{i \leq j} s_i$$

for every index j.

Now suppose no curve σ_j with $1 \leq j \leq d$ is useful. Then we must have $s_j^2 > n_j$ and therefore

$$s_j^2 > \sum_{i \leq j} s_i$$

for all $j \geq 1$. Trivially, $s_1 \geq 1$ unless γ is simple and thus $d = 1$. A straightforward induction argument implies that $s_j \geq (j+1)/2$ and therefore

$$n \;\geq\; n_d \;\geq\; \sum_{i \leq d} \frac{i+1}{2} \;\geq\; \frac{1}{2}\binom{d+2}{2} \;>\; \frac{d^2}{4}.$$

We conclude that $d \leq 2\sqrt{n}$, which completes the proof. ◀

▶ **Theorem 4.4.** *Every closed curve γ in the plane with n vertices can be simplified in $O(n^{3/2})$ homotopy moves.*

Proof. If γ has depth $O(\sqrt{n})$, Lemma 4.1 and the trivial upper bound $D(\gamma) \leq (n+1) \cdot depth(\gamma)$ imply that we can simplify γ in $O(n^{3/2})$ homotopy moves. For purposes of analysis, we charge $O(\sqrt{n})$ of these moves to each vertex of γ.

Otherwise, let σ be an arbitrary useful closed curve chosen according to Lemma 4.3. Suppose the interior tangle of σ has m vertices, s strands, and depth d. Lemma 4.3 implies that $d = O(\sqrt{n})$, and the definition of useful implies that $s \leq \sqrt{m}$, which is $O(\sqrt{n})$. Thus, by Lemma 4.2, we can tighten the interior tangle of σ in $O(md + ms) = O(m\sqrt{n})$ moves. This simplification removes at least $m - s^2/2 \geq m/2$ vertices from γ, as the resulting tight tangle has at most $s^2/2$ vertices. Again, for purposes of analysis, we charge $O(\sqrt{n})$ moves to each deleted vertex. We then recursively simplify the resulting closed curve.

In either case, each vertex of γ is charged $O(\sqrt{n})$ moves as it is deleted. Thus, simplification requires at most $O(n^{3/2})$ homotopy moves in total. ◀

5 Higher-Genus Surfaces

Finally, we consider the natural generalization of our problem to closed curves on orientable surfaces of higher genus. Because these surfaces have non-trivial topology, not every closed curve is homotopic to a single point or even to a simple curve. A closed curve is *contractible* if it is homotopic to a single point. We call a closed curve *tight* if it has the minimum number of self-intersections in its homotopy class.

5.1 Lower Bounds

Although defect was originally defined as an invariant of *planar* curves, Polyak's formula $\mathit{defect}(\gamma) = -2\sum_{x \between y} \operatorname{sgn}(x)\operatorname{sgn}(y)$ extends naturally to closed curves on any orientable surface; homotopy moves change the resulting invariant exactly as described in Figure 2.1. Thus, Lemma 2.1 immediately generalizes to any orientable surface as follows.

▶ **Lemma 5.1.** *Let γ and γ' be arbitrary closed curves that are homotopic on an arbitrary orientable surface. Transforming γ into γ' requires at least $|\mathit{defect}(\gamma) - \mathit{defect}(\gamma')|/2$ homotopy moves.*

The following construction implies a quadratic lower bound for simplifying noncontractible curves on orientable surfaces with any positive genus.

▶ **Lemma 5.2.** *For any positive integer n, there is a closed curve on the torus with n vertices and defect $\Omega(n^2)$ that is homotopic to a simple closed curve but not contractible.*

Proof. Without loss of generality, suppose n is a multiple of 8. The curve γ is illustrated on the left in Figure 5.1. The torus is represented by a rectangle with opposite edges identified. We label three points a, b, c on the vertical edge of the rectangle and decompose the curve into a red path from a to b, a green path from b to c, and a blue path from c to a. The red and blue paths each wind vertically around the torus, first $n/8$ times in one direction, and then $n/8$ times in the opposite direction.

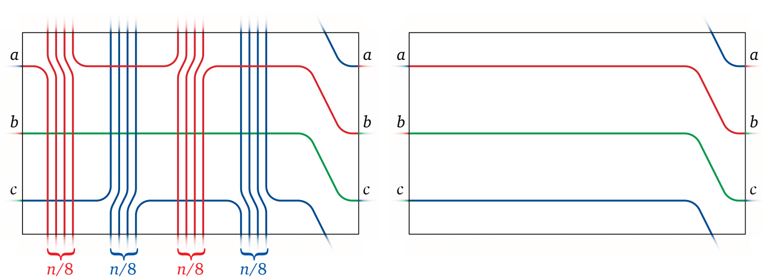

Figure 5.1 A curve γ on the torus with defect $\Omega(n^2)$ and a simple curve homotopic to γ.

As in previous proofs, we compute the defect of γ by describing a sequence of homotopy moves that simplify the curve, while carefully tracking the changes in the defect that these moves incur. We can unwind one turn of the red path by performing one 2→0 move, followed by $n/8$ 3→3 moves, followed by one 2→0 move, as illustrated in Figure 5.2. Repeating this sequence of homotopy moves $n/8$ times removes all intersections between the red and green paths, after which a sequence of $n/4$ 2→0 moves straightens the blue path, yielding the simple curve shown on the right in Figure 5.1. Altogether, we perform $n^2/64 + n/2$ homotopy moves, where each 3→3 move increases the defect of the curve by 2 and each 2→0 move decreases the defect of the curve by 2. We conclude that $\mathit{defect}(\gamma) = -n^2/32 + n$. ◀

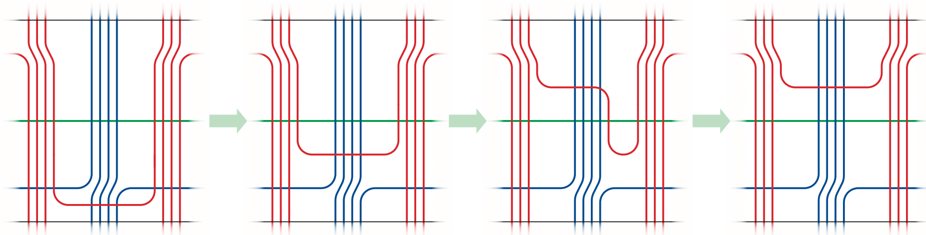

Figure 5.2 Unwinding one turn of the red path.

▶ **Theorem 5.3.** *Simplifying a closed curve with n crossings on a torus requires $\Omega(n^2)$ homotopy moves in the worst case, even if the curve is homotopic to a simple curve.*

5.2 Upper Bounds

Hass and Scott proved that any non-simple closed curve on any orientable surface that is homotopic to a simple closed curve contains either a simple (in fact empty) contractible loop or a simple contractible lens [24, Theorem 1]. It follows immediately that any such curve can be simplified in $O(n^2)$ moves using Steinitz's algorithm; Theorem 5.3 implies that the upper bound is tight for non-contractible curves.

For the most general setting, where the given curve is not necessarily homotopic to a simple closed curve, we are not even aware of a *polynomial* upper bound! Steinitz's algorithm does not work here; there are curves with excess self-intersections but no simple contractible loops or lenses [24]. Hass and Scott [25] and De Graff and Schrijver [14] independently proved that any closed curve on any surface can be simplified using a *finite* number of homotopy moves that never increase the number of self-intersections. Both proofs use discrete variants of curve-shortening flow; for sufficiently well-behaved curves and surfaces, results of Grayson [22] and Angenent [5] imply a similar result for differential curvature flow. Unfortunately, without further assumptions about the precise geometries of both the curve and the underlying surface, the number of homotopy moves cannot be bounded by any function of the number of crossings; even in the plane, there are closed curves with three crossings for which curve-shortening flow alternates between a 3→3 move and its inverse arbitrarily many times. Paterson [34] describes a combinatorial algorithm to compute a simplifying sequence of homotopy moves without such reversals, but she offers no analysis of her algorithm.

We conjecture that any contractible curve on any surface can be simplified with at most $O(n^{3/2})$ homotopy moves, and that arbitrary curves on any surface can be simplified with at most $O(n^2)$ homotopy moves.

References

1 Francesca Aicardi. Tree-like curves. In Vladimir I. Arnold, editor, *Singularities and Bifurcations*, volume 21 of *Advances in Soviet Mathematics*, pages 1–31. Amer. Math. Soc., 1994.

2 Sheldon B. Akers, Jr. The use of wye-delta transformations in network simplification. *Oper. Res.*, 8(3):311–323, 1960.

3 James W. Alexander. Combinatorial analysis situs. *Trans. Amer. Math. Soc.*, 28(2):301–326, 1926.

4 James W. Alexander and G. B. Briggs. On types of knotted curves. *Ann. Math.*, 28(1/4):562–586, 1926–1927.

5 Sigurd Angenent. Prabolic equations for curves on surfaces: Part II. Intersections, blow-up and generalized solutions. *Ann. Math.*, 133(1):171–215, 1991.

6 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.

7 Vladimir I. Arnold. Plane curves, their invariants, perestroikas and classifications. In Vladimir I. Arnold, editor, *Singularities and Bifurcations*, volume 21 of *Adv. Soviet Math.*, pages 33–91. Amer. Math. Soc., 1994.

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

9 Hsien-Chih Chang and Jeff Erickson. Electrical reduction, homtoopy moves, and defect. Preprint, October 2015. URL: http://arxiv.org/abs/1510.00571.

10 Sergei Chmutov, Sergei Duzhin, and Jacob Mostovoy. *Introduction to Vassiliev knot invariants*. Cambridge Univ. Press, 2012.

11 Charles J. Colbourn, J. Scott Provan, and Dirk Vertigan. A new approach to solving three combinatorial enumeration problems on planar graphs. *Discrete Appl. Math.*, 60:119–129, 1995.

12 Yves Colin de Verdière, Isidoro Gitler, and Dirk Vertigan. Réseaux électriques planaires II. *Comment. Math. Helvetici*, 71(144–167), 1996.

13 John H. Conway. An enumeration of knots and links, and some of their algebraic properties. In by:John Leech, editor, *Computational Problems in Abstract Algebra*, pages 329–358. Pergamon Press, 1970.

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

15 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.

16 Chaim Even-Zohar, Joel Hass, Nati Linial, and Tahl Nowik. Invariants of random knots and links. Preprint, November 2014. URL: http://arxiv.org/abs/1411.3308.

17 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.

18 Thomas Aurelio 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.* PhD thesis, Univ. California Berkeley, 1985.

19 George K. Francis. The folded ribbon theorem: A contribution to the study of immersed circles. *Trans. Amer. Math. Soc.*, 141:271–303, 1969.

20 Carl Friedrich Gauß. Nachlass. I. Zur Geometria situs. In *Werke*, volume 8, pages 271–281. Teubner, 1900. Originally written between 1823 and 1840.

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

22 Matthew A. Grayson. Shortening embedded curves. *Ann. Math.*, 129(1):71–111, 1989.
23 Branko Grünbaum. *Convex Polytopes.* Number XVI in Monographs in Pure and Applied Mathematics. John Wiley & Sons, 1967.
24 Joel Hass and Peter Scott. Intersections of curves on surfaces. *Israel J. Math.*, 51:90–120, 1985.
25 Joel Hass and Peter Scott. Shortening curves on surfaces. *Topology*, 33(1):25–43, 1994.
26 Chuichiro Hayashi and Miwa Hayashi. Minimal sequences of Reidemeister moves on diagrams of torus knots. *Proc. Amer. Math. Soc.*, 139:2605–2614, 2011.
27 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.
28 Noburo Ito and Yusuke Takimura. (1,2) and weak (1,3) homotopies on knot projections. *J. Knot Theory Ramif.*, 22(14):1350085 (14 pages), 2013. Addendum in *J. Knot Theory Ramif.* 23(8):1491001 (2 pages), 2014.
29 Arthur Edwin Kennelly. Equivalence of triangles and three-pointed stars in conducting networks. *Electrical World and Engineer*, 34(12):413–414, 1899.
30 Mikhail Khovanov. Doodle groups. *Trans. Amer. Math. Soc.*, 349(6):2297–2315, 1997.
31 Hiroyuki Nakahara and Hiromitsu Takahashi. An algorithm for the solution of a linear system by Δ-Y transformations. *IEICE Trans. Fundamentals of Electronics, Communications and Computer Sciences*, E79-A(7):1079–1088, 1996.
32 Steven D. Noble and Dominic J. A. Welsh. Knot graphs. *J. Graph Theory*, 34(1):100–111, 2000.
33 Tahl Nowik. Complexity of planar and spherical curves. *Duke J. Math.*, 148(1):107–118, 2009.
34 Jane M. Paterson. A combinatorial algorithm for immersed loops in surfaces. *Topology Appl.*, 123:205–234, 2002.
35 Michael Polyak. Invariants of curves and fronts via Gauss diagrams. *Topology*, 37(5):989–1009, 1998.
36 Kurt Reidemeister. Elementare Begründung der Knotentheorie. *Abh. Math. Sem. Hamburg*, 5:24–32, 1927.
37 Neil Robertson and Paul D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Comb. Theory Ser. B*, 52(2):153–190, 1991.
38 Neil Robertson, Paul D. Seymour, and Robin Thomas. Quickly excluding a planar graph. *J. Comb. Theory Ser. B*, 62(2):232–348, 1994.
39 Alexander Russell. The method of duality. In *A Treatise on the Theory of Alternating Currents*, chapter XVII, pages 380–399. Cambridge Univ. Press, 1904.
40 Xiaohuan Song. Implementation issues for Feo and Provan's delta-wye-delta reduction algorithm. M.Sc. Thesis, University of Victoria, 2001.
41 Ernst Steinitz. Polyeder und Raumeinteilungen. *Enzyklopädie der mathematischen Wissenschaften mit Einschluss ihrer Anwendungen*, III.AB(12):1–139, 1916.
42 Ernst Steinitz and Hans Rademacher. *Vorlesungen über die Theorie der Polyeder: unter Einschluß der Elemente der Topologie*, volume 41 of *Grundlehren der mathematischen Wissenschaften.* Springer-Verlag, 1934. Reprinted 1976.
43 Klaus Truemper. On the delta-wye reduction for planar graphs. *J. Graph Theory*, 13(2):141–148, 1989.
44 Leslie S. Valiant. Universality considerations in VLSI circuits. *IEEE Trans. Comput.*, C-30(2):135–140, 1981.
45 Gert Vegter. Kink-free deformation of polygons. In *Proceedings of the 5th Annual Symposium on Computational Geometry*, pages 61–68, 1989.