Highly Non-Convex Crossing Sequences

by

Andrew McConvey

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Combinatorics and Optimization

Waterloo, Ontario, Canada, 2012

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

For a given graph, G, the crossing number $\mathrm{cr}_k(G)$ denotes the minimum number of edge crossings when a graph is drawn on an orientable surface of genus k. The sequence $\mathrm{cr}_0(G), \mathrm{cr}_1(G), ...$ is said to be the crossing sequence of a G. An equivalent definition exists for non-orientable surfaces.

In 1983, Jozef Širáň proved that for every decreasing, convex sequence of non-negative integers, there is a graph G such that this sequence is the crossing sequence of G. This main result of this thesis proves the existence of a graph with non-convex crossing sequence of arbitrary length.

Acknowledgements

The research done for this thesis was in collaboration with Bruce Richter and Gelasio Salazar. Without the numerous meetings and discussions with them, this work would not have been possible. I would also like to give special thanks to Bruce for all his suggestions and corrections.

Dedication

For Jian and Jeannie. Thank you for always going out of your way to help.

Table of Contents

List of Figures

Chapter 1

Introduction

One of the most famous results in Topological Graph Theory is Kuratowski's Theorem, which characterizes the set of graphs that can be drawn on the plane with no edge crossings. The theorem states that a graph is planar if and only if none of its subgraphs is a subdivision of K_5 or $K_{3,3}$.

A drawing of a graph on a surface Σ is a mapping of the vertices of G to distinct points on the surface and the edges to simple curves. For each edge $\{u, v\} \in E$, the points corresponding to u and v will be the endpoints of the curve corresponding to $\{u, v\}$. We specify that at most two curves can cross at a single point and that no curve can intersect a vertex except at its endpoints.

If a graph G is not planar, then two natural questions arise: what is the crossing number of G and what is the genus of G? Both of these questions have been well studied and, in general, are hard problems to solve.

A *crossing* is an intersection of two edges that is not a vertex. The crossing number of a graph, $\mathrm{cr}(G)$, is the minimum number of crossings over all possible drawings on the plane. The first problem in the field was posed by Turán in 1944 and remains open. The problem is to determine the crossing number of the complete bipartite graph, $\mathrm{cr}(K_{m,n})$.

A closed surface is characterized topologically by its orientability and its genus. For an orientable surface, the genus is simply the number of handles on the surface. If the surface is non-orientable, then the genus is the number of crosscaps.

Given a drawing of a graph on the plane with a crossing, one can introduce either a handle or a crosscap to the surface in order to remove the crossing. The result is a drawing on a surface of genus 1 with one fewer crossing. This process can be repeated until the graph is embedded in a surface, meaning that it is drawn with no crossings.

The *genus* of the graph is smallest positive integer g such that the graph can be embedded on the surface of genus g. A graph has an orientable and non-orientable genus. The orientable and non-orientable genus are not necessarily equal and both are at most $\mathrm{cr}(G)$.

Despite these questions having been well studied, it is interesting that very little work has been done on the natural combination of the two: crossing numbers on more general surfaces. The k^{th} orientable crossing number of a graph is defined to be the minimum number of crossings over all possible drawings of G on the orientable surface of genus k and is denoted $\mathrm{cr}_k(G)$. There is an analogous definition for non-orientable surfaces and is denoted $\widetilde{\mathrm{cr}}_k(G)$.

The *orientable crossing sequence* of a graph is defined to be the sequence $\mathrm{cr}_0(G)$, $\mathrm{cr}_1(G), \ldots, \mathrm{cr}_g(G) = 0$, where g is the genus of G. The non-orientable case is also defined in the same fashion. We will generally refer to the orientable version simply as the crossing sequence unless otherwise specified.

As we described above, we can remove at least one crossing by adding a single handle or crosscap to the surface. Therefore, we see that the crossing sequence is a finite sequence of positive integers strictly decreasing to zero.

To date, there have only been 3 major results in the field of crossing sequences. The first was by Širáň [6] in 1983. He introduced the notion of the crossing sequence and showed that for any convex sequence there is a graph with that sequence as its crossing sequence. A decreasing sequence $a_0, a_1, \ldots$ is *convex* if and only if $a_{j-1} - a_j \geq a_j - a_{j+1}$, for any j. He conjectured that every crossing sequence is convex. If true, this conjecture would have completely characterized the set of crossing sequences.

In 2001, Archdeacon, Bonnington, and Širáň [1] found a counter example to Širáň's conjecture. For any positive integers $a > b$, they proved the existence of a graph with non-orientable crossing sequence $a, b, 0$. In particular, when $a - b < b$, this sequence is non-convex. Additionally, they found a counter-example for the orientable case which also had length 3. However, this example lacked the full generality of its non-orientable counterpart.

Finally, in 2011, DeVos, Mohar, and Šámal [3] proved a similar result for the orientable case. Given two positive integers $a > b$, they proved the existence of a graph with orientable crossing sequence $a, b, 0$.

The main result of this paper proves the existence of a graph with non-convex crossing sequence of arbitrary length. In particular, given any $N \in \mathbb{N}$ and $\varepsilon > 0$, there exists a graph such that $(\mathrm{cr}_{n-1} - \mathrm{cr}_n)/(\mathrm{cr}_n - \mathrm{cr}_{n+1}) < \varepsilon$ for some $n \geq N$.

In Chapter 2, we will describe several techniques and gadgets which will allow for easier analysis of the crossing number. Chapter 3 provides several lemmas, including the first result of this thesis, that will used extensively in our analysis. Chapter 4 contains a thorough review of the previous results described above. Finally, in Chapter 5, we describe our graphs $G_{n,k}$ and prove our main result:

Theorem 1.0.1. *Let n be a positive integer and $k \geq 10$ be an even positive integer. Then there exists a graph G that has crossing sequence given by:*

$$\begin{aligned} \mathrm{cr}_0(G) &= 2n^2k^2 \\ \mathrm{cr}_{2g-1}(G) &= [nk-(g-1)k-1]nk \quad \text{for } 1 \leq g \leq n \\ \mathrm{cr}_{2g}(G) &= [nk-(g-1)k-2]nk \quad \text{for } 1 \leq g \leq n \\ \mathrm{cr}_{2n+1}(G) &= 0 \end{aligned}$$

Note that, for $1 \leq g < n$,

$$\begin{aligned} \mathrm{cr}_{2g-1}(G_{n,k}) - \mathrm{cr}_{2g}(G_{n,k}) &= nk \\ \mathrm{cr}_{2g}(G_{n,k}) - \mathrm{cr}_{2g+1}(G_{n,k}) &= nk(k-1)\,. \end{aligned}$$

In particular, for such g, the three term sequence $\mathrm{cr}_{2g-1}(G_{n,k}), \mathrm{cr}_{2g}(G_{n,k}), \mathrm{cr}_{2g+1}(G_{n,k})$ is not convex.

Chapter 2

Techniques and Gadgets

2.1 Cutting a Surface

Much of our analysis of $G_{n,k}$ will center around the homotopy classes of our surface. For completeness, we recall the definition of homotopy.

Let c_1 and c_2 be continuous maps from X to Y. We say that c_1 and c_2 are *homotopic* if there exists a continuous map $H : [0,1] \times X \to Y$ such that $H(0,x) = c_1(x)$ and $H(1,x) = c_2(x)$ for every $x \in X$.

In particular, we will often focus on the homotopy of cycles in the graph. We say a cycle is *contractible* if it is homotopic to a point. On the plane, every cycle is contractible. If a cycle can not be continuously deformed to a point, we say it is *non-contractible*.

A *separating* cycle divides the surface into two disjoint regions, one on each side of the curve. Every contractible curve is separating, but the converse is not true. Similarly, a *non-separating* curve does not divide the surface. Such a curve must be non-contractible. Figure 2.1 gives an example of two non-contractible curves.

Suppose we have drawing $\mathcal{D}$ of a graph G on a given surface Σ of genus g. Suppose G contains a cycle C that contains no crossings and is non-contractible in $\mathcal{D}$. Cutting Σ along C leaves a new surface Σ' with two holes, each bounded by a copy of C. We will call these copies C_1 and C_2. Cap each of the holes with a disc. There are two cases.

Case I (C is Non-Separating) If C is non-separating, Σ' is now a closed surface with genus $g-1$. Since C was uncrossed in $\mathcal{D}$, the drawing extends to a drawing $\mathcal{D}'$ in

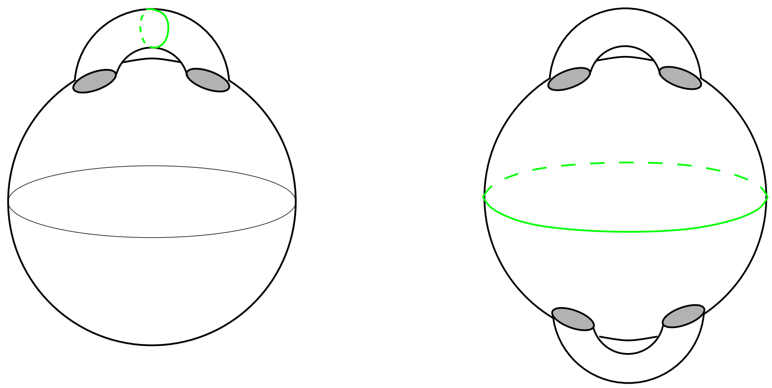

Figure 2.1: Two non-contractible curves. The curve on the left is non-separating and the curve on the right is separating.

Σ'. In this drawing, there are two copies of C and every edge and vertex in $G \setminus C$ is found exactly once.

Suppose C' is another cycle in G and assume that it is not self-crossing in $\mathcal{D}$. If C' is disjoint from C, then it is drawn as a closed curve in $\mathcal{D}'$. However, if it is not disjoint from C, this may not be the case.

The case that C and C' intersect at just one point s is an important one in this thesis. If this point is not a tangent point in $\mathcal{D}$ (i.e. C' crosses C), then C' will not be a simple closed curve in $\mathcal{D}'$. Since there are two copies of C in $\mathcal{D}'$, each of these contains a copy of the point s, which we will call s_1 and s_2 respectively. In $\mathcal{D}'$, the cycle C' appears as an arc on Σ' with end points s_1 and s_2.

More generally, we can distinguish the two sides of the curve C in our original drawing. Any edge incident to one side of the curve in $\mathcal{D}$, say the left, will be incident to C_1 in $\mathcal{D}'$. Any edge incident to the other side of the curve in $\mathcal{D}$, now the right, will be incident to C_2 in $\mathcal{D}'$.

Case II (*C* is Separating) If C is a separating curve, cutting and capping Σ along C leaves two disjoint surfaces, Σ_1 and Σ_2 with genus g_1 and g_2, respectively. Note that $g = g_1 + g_2$.

The drawing $\mathcal{D}$ induces drawings of two subgraphs of G, namely $G_1 := G \cap \Sigma_1$ in Σ_1 and $G_2 := G \cap \Sigma_2$ in Σ_2. Note that $G_1 \cup G_2 = G$ and $G_1 \cap G_2 = C$.

Let C_1 be the copy of C in G_1 and C_2 the copy of C in G_2. The cycles C_1 and C_2 each bound a disc which is also a face on Σ_1 and Σ_2 respectively. Since C was uncrossed in $\mathcal{D}$ every remaining edge and vertex must be drawn on exactly one of Σ_1 or Σ_2.

As in the non-separating case, we can distinguish the two sides of the curve C in our original drawing. If an edge e was incident to one side of the curve in $\mathcal{D}$, say the left, it will be found in G_1 and incident to C_1. This also implies that e will be contained on the surface Σ_1. Any edge incident to the other side of the curve in $\mathcal{D}$, now the right, will be incident to C_2 and be drawn on Σ_2. No edge could be on both Σ_1 and Σ_2 because that would imply that it was in both regions of Σ that were bounded by C.

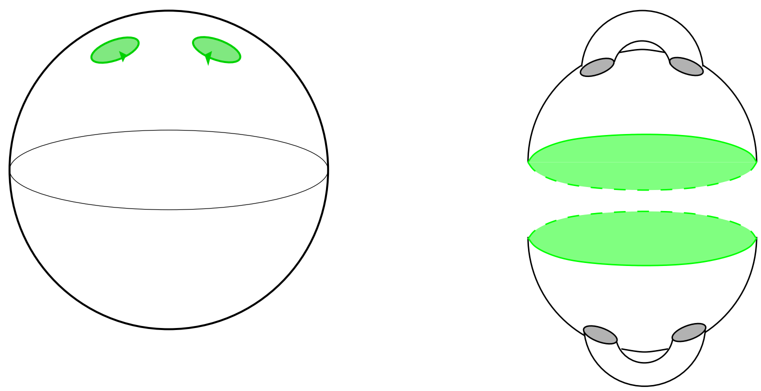

Figure 2.2: The surfaces of Figure 2.1 after cutting and capping.

We will call this process "cutting and capping" along C. It is particularly useful because the non-contractibility of C implies each new surface that is created has genus strictly less than g. This will allow us to use induction on the genus of the surface.

2.2 Weighted Edges

Weighted edges are often used in crossing numbers to dictate which edges can and cannot be crossed. Suppose we have a graph in which an edge e has weight $w_e \in \mathbb{N}$. If edges e and e' cross in a drawing of G, we say that the *weight of the crossing* is $w_e \cdot w_{e'}$.

The *weighted crossing number* $\mathrm{wcr}_g(G)$ of a graph G is defined to be the minimum sum of weighted crossings in any drawing of the graph in the surface of genus g.

It is shown in [1] that, for any weighted graph G' with positive integer weights, there is an unweighted graph G with $\mathrm{wcr}_g(G') = \mathrm{cr}_g(G)$, for all $g \geq 0$. Given such a weighted graph G', we will describe a recursive process to find the unweighted graph G with the same crossing number.

Let $e \in E(G')$ be an edge such that $w_e > 1$. Define a new weighted graph G'' by replacing e with w_e parallel edges of weight 1. Clearly G'' has fewer edges of weight greater than 1 and we will show that $\text{wcr}_g(G'') = \text{wcr}_g(G')$

Consider an optimal weighted drawing $\mathcal{D}'$ of G' on a surface of genus g. Obtain the drawing $\mathcal{D}''$ of G'' by removing the curve associated to edge e from $\mathcal{D}'$ and replacing it with w_e parallel curves, all drawn near to each other in the surface. Each of these edges has weight one.

Since each parallel edge in $\mathcal{D}''$ crosses the exact same weighted edges as e in $\mathcal{D}'$, the weighted crossing number of $\mathcal{D}''$ must the same as the weighted crossing number of $\mathcal{D}'$. In particular, we know that $\text{wcr}_g(G') \geq \text{wcr}_g(G'')$.

Now suppose we have an optimal drawing of G'' on a surface of genus g. Take the parallel edge that contributes the least to the weighted crossing number. This drawing can be modified so that all w_e of the edges run alongside this minimal edge without increasing the crossing number of the drawing.

Replacing these now parallel edges with one edge of weight w_e gives a drawing of G' with the same weighted crossing number. So $\text{wcr}_g(G'') \geq \text{wcr}_g(G')$ and, from the previous paragraph, equality must hold. Note that this is true for any genus g.

So given a weighted graph G', we have found a new graph G'' with the same weighted crossing number and fewer edges with weight greater than 1. The process can be repeated until we have a weighted graph G'' in which all edges have weight 1 and $\text{wcr}_g(G'') = \text{wcr}_g(G')$.

From G'', we wish to obtain a simple graph G. First, we subdivide any parallel edges and note that this does not affect the weighted crossing number. Now, G'' is a weighted graph with no parallel edges and each edge has weight one. Let G be the simple, unweighted graph with vertex set and edge set equivalent to the vertex set and edge set of G''. Since each edge in G'' has weight 1, it is clear that $\text{cr}_g(G) = \text{wcr}_g(G'') = \text{wcr}_g(G')$.

Lastly, we introduce the similar concept of *thick* edges. A thick edge is simply an edge with infinite weight. The weighted crossing number is finite if and only if the subgraph of G consisting of all the thick edges is planar. In this case, thick edges can be treated simply as weighted edges with weight greater than $\text{cr}_0(G)$.

2.3 Rigid Vertices

Suppose there is a graph with vertex v and incident to edges $e_1, e_2, \ldots, e_n$. If, in the given drawing and proceeding clockwise from e_1, we find the order of the incident edges to be $e_1, e_2, \ldots, e_n$, then we say that the *(clockwise) rotation* of v is $\pi(v) := e_1, \ldots, e_n$. We may similarly define the *counter-clockwise rotation*.

Rigid vertices were used by DeVos et al. [3] as a means to prescribe the rotations of a vertex. For completeness, we will include their discussion here. A *rigid vertex* is a vertex v whose incident edges have a local cyclic rotation π_v, prescribed up to inversion. This means, in any drawing of the graph, either the clockwise rotation or the counter-clockwise rotation of v is given by $\pi(v)$.

We show that for any graph with rigid vertices with crossing number cr_g, there exists a graph without rigid vertices with the same crossing number.

Suppose G is a connected graph with rigid vertex v. We will create the new graph G' by replacing v with a copy of graph $V_{n,\deg(v)}$. Suppose the prescribed rotation of v is $e_1, e_2, \ldots, e_{\deg(v)}$. To obtain G' from G, subdivide each e_i n times, creating new vertices $v_{i,1}, v_{i,2}, \ldots, v_{i,n}$, in order from v. Then add thick cycles (i.e cycles with thick edges) $c_j := v_{1,j}, v_{2,j}, \ldots, v_{\deg(v),j}, v_{1,j}$, for each $j \in \{1, 2, \ldots, n\}$. We will show that, when $n \geq 3g + 2$, $\mathrm{cr}_g(G) = \mathrm{cr}_g(G')$.

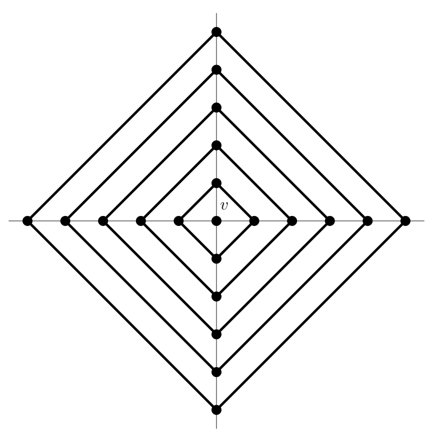

Figure 2.3: A drawing of $V_{n,\deg(v)}$ (here $n = 5$ and $\deg(v) = 4$).

Suppose we have an optimal drawing of G on a surface of genus g. Because v is rigid, we can replace v with a drawing of $V_{n,\deg(v)}$ by the process described above. Since all new edges can be placed within an ε-neighborhood of v, the process will add no more crossings. Therefore, $\mathrm{cr}_g(G) \geq \mathrm{cr}_g(G')$.

If there is an optimal drawing of G' such that $V_{n,\deg(v)}$ is drawn as in Figure 2.3 (i.e. with the thick cycles c_i of $V_{n,\deg(v)}$ contractible and nested), then we can replace $V_{n,\deg(v)}$ with a rigid vertex v without introducing additional crossings. In this case, $\mathrm{cr}_g(G) \leq \mathrm{cr}_g(G')$ and our proof is complete.

Suppose we have an optimal drawing $\mathcal{D}$ of G'. Following the argument of DeVos et al., we will show that there is another optimal drawing $\mathcal{D}'$ of G' so that $V_{n,\deg(v)}$ is drawn as in Figure 2.3. Consider the n disjoint cycles $c_1, \ldots, c_n$ and note that no thick cycle can be crossed.

Two of the c_i's are Contractible Suppose c_i and c_j are contractible with $i < j$ in $\mathcal{D}$. Let C_i be the disc bounded by c_i. Note that there are exactly two connected components of $G' \setminus c_i$, one containing c_j and the other containing v.

Since c_i is thick, each component must be completely contained in a face bounded by c_i and, since C_i is a disc, we can assume that at least one component of the graph is contained in the other face. The component containing v is planar, so we conclude that it is contained in C_i and that the other component is contained in the other face.

Now, we can redraw all the cycles c_k to be in an ε-neighborhood of c_i, so that each C_k is contained in C_{k+1}. Drawing the cycles in this way introduces no new crossings to the graph and creates a drawing of $V_{n,\deg(v)}$ as in Figure 2.3.

No two of the c_i's are Contractible There can be at most $3g-3$ disjoint, non-contractible, pairwise non-homotopic cycles on the surface of genus g [4, Prop. 4.2.6]. Since $n \geq 3g+2$, there must exist two homotopic cycles c_i and c_j with $i < j < n$.

Since c_i and c_j are thick, each component of $G' \setminus (c_i \cup c_j)$ must be contained in a face bounded by $c_i \cup c_j$. Let K be the connected component containing c_n. Each other connected component is planar and can either be drawn in the cylinder bounded by c_i and c_j or within a ε-neighborhood of c_i. Thus, we can assume that K is not contained in the cylinder bounded by c_i and c_j. Moreover, K is incident to only one side of c_j.

Cut and cap along c_j. If c_j was separating, then consider the new surface Σ_1 with genus $g_1 < g$ containing $K \cup c_j$. If c_j was non-separating, then the new surface Σ' has genus $g-1$. Delete the copy of c_j not incident to K and all vertices and edges in $G' \setminus (K \cup c_j)$. In both cases, we are left with a drawing of $K \cup c_j$ with no more crossings on a surface of strictly smaller genus. Moreover, in this new drawing, c_j is contractible and uncrossed.

Since $V_{j,\deg(v)} = G' \setminus K$ is planar, it can be embedded in the disc bounded by c_j. We must also redraw the cycles $c_{j+1}, \ldots, c_n$ alongside c_j. If we do so, we obtain a drawing of G' in Σ' with no more crossings than the original drawing. Since the genus of Σ' is strictly smaller than g, this contradicts the optimality of $\mathcal{D}$, so no such drawing could exist.

We can therefore conclude that $\mathrm{cr}_g(G) = \mathrm{cr}_g(G')$. So for any graph G with rigid vertices, there exists a graph with no rigid vertex and the same crossing sequence.

Note that the proof would still hold if we replaced G with a weighted graph with rigid vertices and instead show that there exists a new weighted graph with no rigid vertex and the same weighted crossing sequence.

Chapter 3

Disjoint Cycles and Bouquets

In this section, we introduce two theorems which we will use to analyze the crossing sequence of our graph. The various cycles contained in our graph will often help us to obtain a bound on the crossing number. The two theorems limit the ways in which the cycles can be drawn on a surface.

3.1 Disjoint Cycles

The following theorem gives one relatively simple, but very useful property of disjoint curves on orientable surfaces. Note that two curves are disjoint on a surface if they do not intersect.

Theorem 3.1.1. *Let Σ be an orientable surface of genus g, and let $\mathcal{C}$ be a set of pairwise disjoint simple closed curves in Σ such that $|\mathcal{C}| > g$. Then some subset of curves in $\mathcal{C}$ separates Σ.*

Theorem 3.1.1 is a fundamental property of orientable surfaces and some sources (e.g. [5]) even define the genus as the maximum number of disjoint closed curves that can be drawn without separating the surface. We provide a short proof using the cutting and capping technique.

Proof. If Σ is a sphere, then every curve in $\mathcal{C}$ is contractible, and hence separating. Since $\mathcal{C}$ contains at least one such curve, we are done. For the case when $g \geq 1$, we proceed by induction.

Pick a curve, $c_1 \in \mathcal{C}$. If c_1 is separating, we are done. Otherwise, cut and cap along it to obtain a surface Σ' of genus $g-1$.

Now $\mathcal{C} \setminus \{c_1\}$ is a set of at least g pairwise disjoint simple closed curves in Σ'. By the induction hypothesis, some subset $C \subseteq \mathcal{C} \setminus \{c_1\}$ separates Σ'.

If the two copies of c_1 are in the same component of $\Sigma \setminus \bigcup_{c' \in C} c'$, then C separates Σ. Otherwise, $C \cup \{c_1\}$ separates Σ. □

3.2 The RBG-Bouquet

When we consider a drawing of the graph $G_{n,k}$ in some surface, we will need to show that, for paths that produce crossings, there are only limited possibilities for how these paths can be drawn. Our next result, the first new result of this thesis, describes in sufficient detail these limitations.

A *bouquet* is a graph with a single vertex. Note that every edge in a bouquet is a loop. Suppose we have a bouquet with the vertex s and ℓ non-contractible loops. These ℓ loops are partitioned into 3 classes, $R := r_1, \ldots, r_i$, $G := g_1, \ldots, g_j$, and $B := b_1, \ldots, b_k$. Each of the ℓ loops is incident to s exactly twice and the rotation of the loops about s is prescribed.

In the first half of the rotation, we see the elements of R followed by the elements of $G \cup B$. Within each their individual class, the elements of R are ascending and the elements of B and G are each descending.

Next, in the second half of the rotation, we see the elements of $B \cup R$ followed by G. This time, within their respective classes the elements of R and B are each ascending and the elements of G are descending. The rotation is therefore $R \to B \cup G \to B \cup R \to G$. Such an embedded bouquet is an *RBG-bouquet*.

We will see that bouquets with this type of rotation can often be created by contracting a path to a point. A small example of the rotation can be seen in Figure 3.1.

It will often be the case that if two paths are homotopic in a drawing of $G_{n,k}$, then they must contain a crossing. By bounding the number of non-homotopic paths, we will be able to find a bound for the crossing number of the graph.

Theorem 3.2.1. *Suppose an RBG-bouquet with vertex s and ℓ loops is embedded with no crossings in the sphere with g handles. If all loops are non-contractible and pairwise non-homotopic, then $\ell \leq 3g$.*

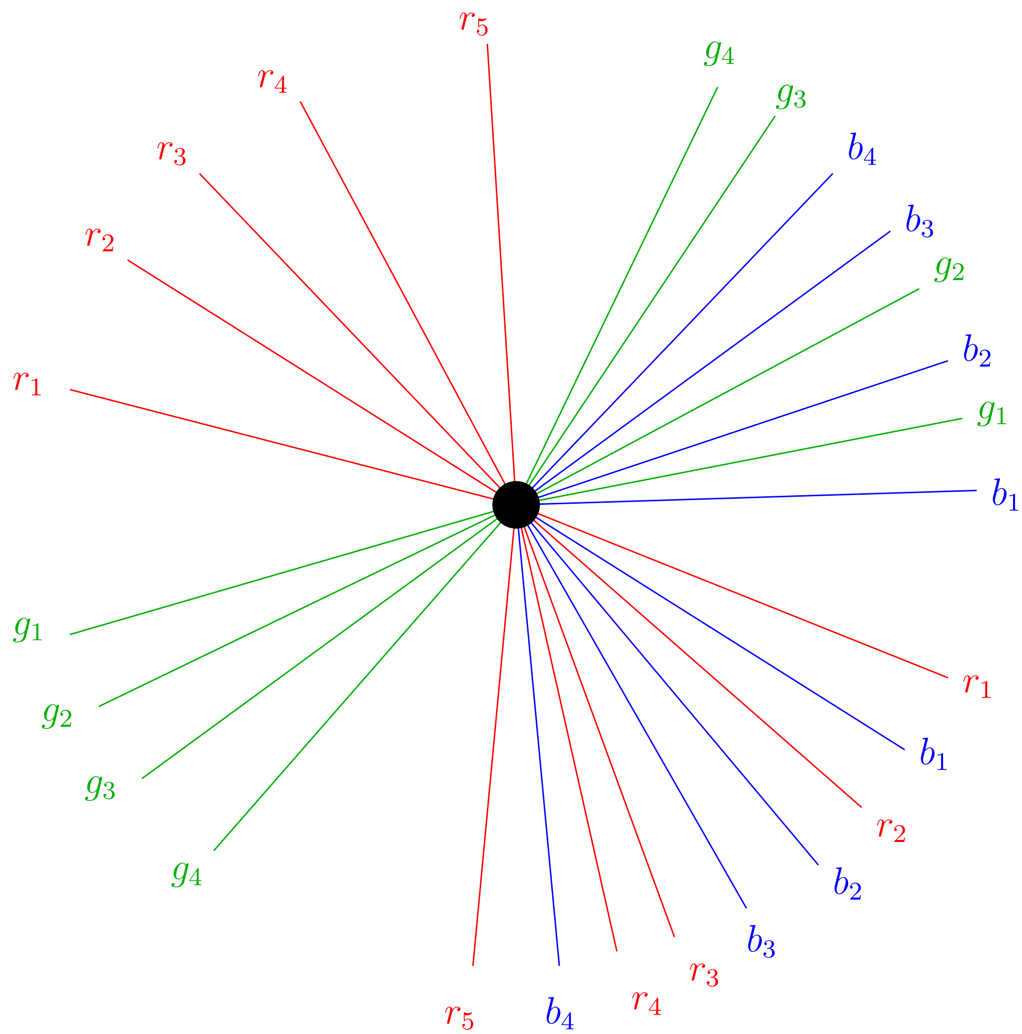

Figure 3.1: Bouquet

If we constrain the rotation of the bouquet even further, then we are able to find a tighter bound. In particular, if all the loops are in B, the following corollary can be proved without significantly modifying the argument.

Corollary 3.2.2. *Suppose an RBG-bouquet with ℓ loops is embedded with no crossings in the sphere with g handles. If $R \cup G = \emptyset$ and all loops are non-contractible and pairwise non-homotopic, then $\ell \leq 3g - 2$.*

We will prove the theorem by induction on the genus g of the surface. We will cut and cap along a given loop and show that almost all of the loops remain non-contractible and pairwise non-homotopic in the new surface. The rotation of s will determine which loop is cut.

Proof. **Base Case**

Suppose $g = 1$. Assume there are four non-contractible, non-homotopic loops e_1, e_2, e_3 and e_4 on the torus.

First note that two cycles are disjoint in a drawing of G are *disjoint* if they have no common points on the surface. Now, restrict the embedding to just two loops, e_i and e_j where $i, j \in \{1, 2, 3, 4\}$. If the rotation of these loops is e_i, e_j, e_j, e_i, then e_j is homotopic to a cycle disjoint from e_i. Since disjoint, non-homotopic, non-contractible cycles on the torus must cross [4, Prop. 4.2.6], the loops e_i and e_j must cross, a contradiction. We conclude that the rotation must be e_i, e_j, e_i, e_j for any choice of i, j. Moreover, we can assume the rotation of all four loops is $e_1, e_2, e_3, e_4, e_1, e_2, e_3, e_4$.

Cut and cap along e_1. From the rotation we know that e_1 was non-separating, so we are left with the sphere Σ' with two copies e_1 and e_1' of e_1, containing the copies s and s', respectively, of s.

The other loops become simple curves with endpoints s and s' in Σ' and any two will separate the sphere into two faces. Consider the faces created by e_2 and e_3. Each face must contain exactly one of e_1 or e_1', as otherwise the two curves would have been homotopic in the original surface.

The loop e_4 must be contained in one of the two regions bounded by e_2 and e_3 and divide the region into two sub-faces. One of the sub-faces is bounded by e_2 and e_4 and a second bounded by e_3 and e_4. Each sub-face must also contain a copy of e_1, as otherwise e_4 would have been homotopic to either e_2 or e_3 in Σ. This cannot occur since there is only one copy of e_1 in the original face.

We conclude that on the torus, there are at most 3 non-contractible, non-homotopic loops in the bouquet and the lemma holds.

The inductive step has two major cases, which depend on the rotation of s.

In the first major case, $B \neq \emptyset$ and the ends of b_1 are consecutive in the rotation of s.

We will cut and cap along b_1, the first loop in B.

Case 1 *Suppose b_1 is separating.* Cutting and capping along b_1 creates two surfaces, each of genus at most $g-1$. Due to the rotation of s, one surface Σ' contains all remaining loops.

None of the remaining loops can be contractible in Σ' as otherwise it would be either contractible in Σ or homotopic to b_1. If two loops are homotopic in Σ', then they bound a cylinder that contains the disc bounded by b_1. If there were two such pairs, then one pair would have been homotopic in Σ.

So after cutting and capping along b_1 and possibly removing one loop from a newly homotopic pair, we have a bouquet with $\ell - 2$ loops on a surface of genus at most $g - 1$. By induction, we have $\ell \leq 3g - 1 \leq 3g$.

Case 2 *Suppose b_1 is non-separating.* We will separate this case into two subcases.

Subcase (i) *Suppose $g = 2$.* Cutting and capping along b_1 yields a torus, Σ'. Recall that there are two copies of b_1 on Σ'. One of the two, which we will continue to call b_1, is still incident to all the other loops of the bouquet. The other copy of b_1, which we will denote b_1', is elsewhere on the surface.

Suppose first that two of the remaining loops, e_1 and e_2, are homotopic on Σ' but were not homotopic on Σ. Together, they bound two faces on Σ', each of which is a cylinder; one cylinder is found to the right of e_1 and the other is found to the left e_1. Since e_1 and e_2 were not homotopic in Σ, each of these cylinders must contain one of b_1 and b_1'. Clearly, e_1 and e_2 can be the only curves with this property.

Additionally, there may be a loop e_3 that is contractible in Σ'. Then the disc that it bounds in Σ' must contain both b_1 and b_1'. To see this, observe that if it contained neither it would have been contractible in Σ and if it contained only one it would have been homotopic to b_1.

The existence of e_3 and the existence of e_1 and e_2 are mutually exclusive as otherwise e_3 must cross the boundary of the cylinder bounded by e_1 and e_2. Removing the newly contractible or one of the homotopic loops, we now have a bouquet with $\ell - 2$ non-contractible, non-homotopic loops on the torus. By induction, $\ell \leq 5 \leq 3g - 1$.

Subcase (ii) *Suppose $g \geq 3$.* In this case, cutting and capping yields a surface Σ' of genus $g - 1$. As in subcase (i), there are 2 copies of b_1 on Σ', each bounding a disc. One of these copies, which we continue to call b_1, is still attached to the bouquet. The other copy, which we call b_1', appears elsewhere on the surface.

Suppose that there are two curves e_1 and e_2 which are homotopic in Σ' but were not homotopic on Σ. The cylinder that they bound must contain either the disc bounded by b_1 or the disc bounded by b_1'. There are two subcases. The first is if this cylinder contains both of the discs and the second is if the cylinder contains exactly one.

Subcase (ii.a) Suppose the cylinder bounded by e_1 and e_2 contains both b_1 and b_1'. Clearly e_1 and e_2 must be the only two such curves, since any other curve with this property must have originally been homotopic to one of them in the original surface.

Suppose there is a third curve e_3 that is contractible in Σ' but was not in Σ. Then the disc it bounds in Σ' must contain both b_1 and b_1'. If it contained neither it would have been contractible in Σ and if it contained just one it would have been homotopic to b_1. Clearly, only one such loop can exist.

Further suppose there is a loop e_3' which is now homotopic to e_1 and e_2, but was homotopic to neither in Σ. Then, the cylinder bounded by e_1 and e_3' must contain either b_1 or b_1', but not both as otherwise e_3' would have been homotopic to either e_1 or e_2 in Σ. Similarly, the cylinder bounded by e_2 and e_3' must contain the other of b_1 and b_1'. Again, only one such e_3' could exist.

The existence of e_3 and e_3' are mutually exclusive, as the disc bounded by e_3 would have to be in distinct cylinders having disjoint interiors. Therefore, there are at least $\ell - 3$ non-homotopic, non-contractible loops in the bouquet on Σ'

Subcase (ii.b) We can assume no two loops bound a cylinder containing both b_1 and b_1', as that case was discussed in (ii.a). Suppose the cylinder bounded by e_1 and e_2 contains b_1 and not b_1'; the other choice is treated in a completely analogous manner. No other curve can be in such a cylinder containing b_1 without being homotopic to either e_1 or e_2 in Σ.

As in our discussion of the double torus, any loop that is contractible in Σ' but not in Σ must contain both b_1 and b_1'. Suppose e_3 is such a contractible curve.

Any cylinder that separates b_1 from b_1' must have its boundary curves within the disc bounded by e_3. This implies that the boundary curves are contractible in Σ, a contradiction.

We conclude that e_1 and e_2 may be newly homotopic in Σ', and there may be other newly homotopic curves whose cylinder bound b_1', but these are the only such examples. Therefore, after cutting and capping along b_1 and removing one loop from two possible newly homotopic pairs, there are at least $\ell - 3$ non-homotopic, non-contractible loops in the bouquet on Σ'.

In both subcases, we have a bouquet of at least $\ell - 3$ loops on a surface of genus $g-1$. By the induction hypothesis, $\ell \leq 3g$, completing the proof.

Corollary 3.2.2: If $R \cup G = \emptyset$, then $\ell \leq 3g - 2$

In this special case, we know s has the rotation $b_1, \ldots, b_\ell, b_\ell, \ldots, b_1$. On the torus, there can only be one loop. If there were two, then b_2 would be homotopic to a cycle disjoint from b_1 and must cross b_1 [4, Prop. 4.2.6]. So when $g = 1$, $\ell \leq 1 = 3g - 2$.

The induction step follows exactly as in Cases 1 and 2 above. However, each time we invoke the induction step, we do so with the hypothesis that $\ell \leq 3g - 2$. We can conclude that when $R \cup G = \emptyset$, $\ell \leq 3g - 2$.

In the second major case, the ends of b_1 are not consecutive.

We now choose the loop along which we will cut and cap. If $B = \emptyset$, the choice of either r_1 or g_1 is arbitrary. If $B \neq \emptyset$, at least one of r_1 and g_1 must be found between the two ends of b_1. We will assume that r_1 is such a loop, although the case for g_1 is completely symmetric, by considering the inverse rotation in which G becomes R.

If r_1 is the only loop, then we are trivially done, so we may assume it is not the only one. Due to the rotation of s, we know that for any other loop x, the ends of r_1 and x occur in the order r_1, x, r_1, x. This implies that every loop is non-separating.

Cutting and capping along r_1 will create a surface Σ' of genus $g-1$, with two copies of r_1 and two copies of s. Let these two copies be r_1 and r_1', containing the point s and s', respectively. Since in the original surface, for each loop x, r_1 and x had rotation r_1, x, r_1, x, each loop now has one end in each copy of s.

Additionally, the original rotation of s can be partitioned into two "halves", with the first half $r_1, y_2, \ldots, y_\ell$, where y_ℓ is either b_1, g_1, or r_ℓ (the last case implies all loops are in R). The second half has rotation $r_1, z_2, \ldots, z_\ell$ with z_2 either r_2, b_1, or g_1.

If z_2 is r_2, then after cutting along r_1, we will contract r_2. Otherwise, we will contract along y_ℓ which is either b_1 or g_1.

If r_2 is contracted, then it must have immediately followed r_1 in both halves of the rotation. Contracting r_2 will create a new bouquet with rotation $y_3, \ldots, y_\ell$ in the first half

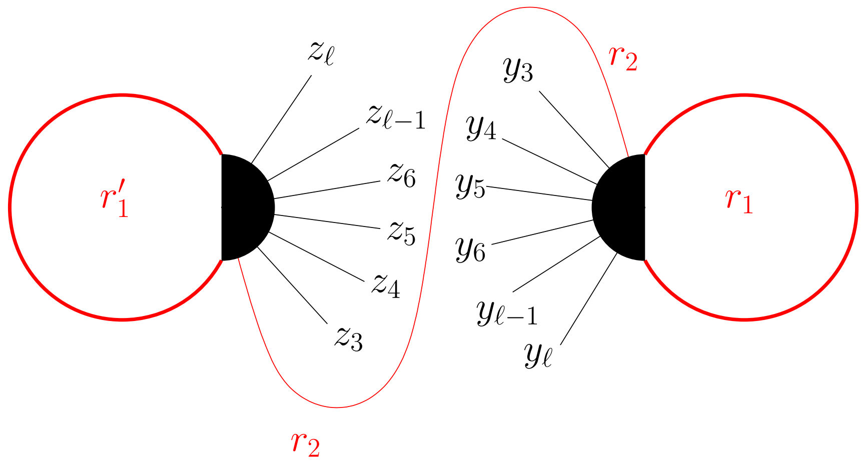

Figure 3.2: The vertices s and s' after cutting along r_1 and before contracting r_2

and $z_3, \ldots, z_\ell$ in the second half. Figure 3.2 shows the bouquet after cutting along r_1 and before contracting r_2.

If g_1 is contracted, then g_1 must be the last loop in both halves of the rotation, so the rotation of the new bouquet is $y_2, \ldots, y_{\ell-1}$ in the first half and $z_2, \ldots, z_{\ell-1}$ in the second half. This is symmetric to the case of contracting r_2.

Finally, if b_1 is contracted, then b_1 was the final loop in the first half of the rotation and immediately followed r_1 in the second half. Contracting b_1 will create a new bouquet with rotation $y_2, \ldots, y_{\ell-1}$ in the first half and $z_3, \ldots, z_\ell$ in the send half. Figure 3.3 shows the bouquet after cutting r_1 and before contracting b_1.

In any case, we are left with a bouquet with $\ell - 2$ loops with the appropriate rotation. We need only determine if there are any newly contractible or newly homotopic loops.

Suppose there are two loops x and y that are newly homotopic. Then they must bound a cylinder in the surface Σ'. Moreover, this cylinder must contain at least one of r_1 and r_1', as otherwise x and y would have been homotopic in Σ.

Let z be the loop that was contracted to re-establish the bouquet. Before contracting z, $x \cup z$ and $y \cup z$ must have bounded the cylinder containing either r_1 or r_1'. Due to the location of r_1 and r_1' in the rotation, the cylinder contains neither or both of r_1 and r_1'. So after contracting z, the cylinder bounded by x and y must contain both r_1 and r_1'.

Clearly, x and y can be the only such newly homotopic loops. If there are no newly contractible loops, deleting y creates a bouquet on Σ' with every loop non-contractible and

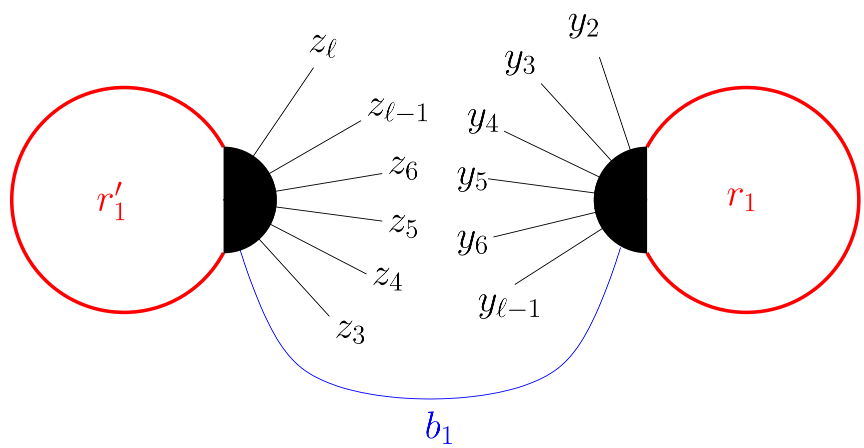

Figure 3.3: The vertices s and s' after cutting along r_1 and before contracting b_1

pairwise non-homotopic. Since we cut and capped along r_1, then contracted another loop, and deleted y, this bouquet has $\ell - 3$ loops. By induction, $\ell \leq 3g$.

We now examine the possibilities when the new bouquet contains newly contractible loops. Since a disc bounded by a newly contractible loop must contain exactly one of r_1 or r_1', the newly contractible loops must be contained in the cylinder bounded by x and y.

Case 1 *Suppose that r_2 is the contracted loop.* If a loop u is newly contractible, then before contraction of r_2, the cycle formed by $r_2 u$ must have bounded a disc. Due to the rotation of s, this disc must contain exactly one of r_1 or r_1'. Moreover, if v is a second such loop, the disc bounded by $r_2 v$ must be disjoint from the disc bounded by $r_2 u$ and contain exactly one of r_1 and r_2'.

Subcase (i) *Suppose no two loops are newly homotopic.* If there is only one newly contractible loop, then after cutting r_1, contracting r_2, and removing the contractible loop, we have a bouquet with $\ell - 3$ non-contractible, non-homotopic loops in a surface of genus $g - 1$. By induction, $\ell \leq 3g$.

Otherwise, suppose that u and v are both newly contractible loops. The order of r_2, u, and v in the first half of the rotation must be r_2, u, v and must be r_2, v, u in the second half. If this were not true, then v would cross the disc bounded by r_2 and u. Any remaining loops must be between u and v in the first half of the rotation and between v and u in the second half, as any other position would imply the loop is in the disc bounded by $r_2 v$ or $r_2 u$.

In particular, any loop between u and v in the first half of the rotation and between v and u the second half must be in B. So, after cutting along r_1, contracting along r_2, and removing u and v, we are left with a drawing with $\ell-4$ loops on the surface of $g-1$ in which all loops are in B. By Corollary 3.2.2 when $R \cup G = \emptyset$, $\ell - 4 \leq 3(g-1) - 2$. So $\ell \leq 3g - 1$.

Subcase (ii) *Suppose x and y are newly homotopic loops.* Since x and y are homotopic and uncrossed, the rotation of x and y must be x, y, y, x. We will assume that x precedes y in the first half of the rotation and y precedes x in the second half.

Let u be a newly contractible loop. The ends of u must be consecutive in Σ' and u must bound a disc containing r_1 or r_1'. Since the cylinder bounded by x and y contains both r_1 and r_1', the cylinder must contain u. Moreover, if v is a second newly contractible loop, it must also be in this cylinder.

We now wish to count the number of loops contained in the cylinder. Recall that r_2 was the loop contracted to reestablish the bouquet, so $x \cup r_2$ and $y \cup r_2$ bounded the cylinder before the contraction of r_2. We will consider the loops before the contraction of r_2.

Any loop contained in the cylinder has one end in s and one end in s'. We will suppose that a loop starts in s and ends in s'. It must start in s either between x and r_2 or between y and r_1. It must end in s' either between r_1' and x or between r_2 and y.

If a loop has an end between x and r_2 at s or between x and r_1' at s', that is equivalent to that loop being incident to the side of the cylinder bounded by $x \cup r_2$. Similarly, if a loop has an end between y and r_1 at s or between r_2 and y at s', this is equivalent to that end being incident to the side of the cylinder bounded by $y \cup r_2$.

Consider two loops z and z' that each have one end in $x \cup r_2$ and one end in $y \cup r_2$. After contracting r_2, these two loops start at the same point on the side of the cylinder bounded by x and end at the same point on the side bounded by y . There can be at most 2 non-homotopic arcs in the cylinder with one end each boundary cycle, so z and z' can be the only such loops.

Now consider the order of loops incident to the side of the cylinder bounded by $x \cup r_2$. Beginning with x at s, we see all loops between x and r_2 at s (that is $x = y_i, y_{i-1}, \ldots, y_2 = r_2$), followed by all loops between r_1' and x at s' (that is $r_1 = z_\ell, \ldots, z_{i'} = x$).

Similarly, if we look at the loops incident to $y \cup r_2$, beginning with y at s, we see all loops between y and r_1 at s (that is $y = y_j, \ldots, y_\ell$), followed by all loops between r_2 and y at s' (that is $r_2 = z_2, \ldots, z_{j'} = y$). Moreover, since the rotation of x and y is x, y, y, x, we have that $i < j$ and $j' < i'$.

If z precedes z' in both of $x \cup r_2$ and $y \cup r_2$, then the columns are homotopic in the cylinder. So z must precede z' in one of the cycles, and z' must precede z in the other. This means that in the original rotation, the rotation of z and z' was z, z', z, z'.

Now we can consider when a loop u has both ends in $x \cup r_2$ or both ends in $y \cup r_2$. In the case where both ends are in $x \cup r_2$, then the loop can either be homotopic to r_2 or homotopic to x. If u is homotopic to x in the cylinder, then it was homotopic to x in the original surface, a contradiction.

If the loop is homotopic to r_2, then it become contractible after contracting r_2 and the disc it bounds contains r_1'. There can only be one such loop and it must have ends y_3 and z_ℓ in the original rotation.

There can be another such loop v which has both ends incident to $y \cup r_2$. The loop v must also be homotopic to r_2 in the cylinder and the disc it bounds must contain r_1. So in the original rotation, the ends of v are y_ℓ and z_3.

Suppose u, v, z, and z' are all loops in the cylinder. Then, the rotation of these four loops must have been $u, z, z'v, v, z, z', u$ in the original rotation of s. However, this is incompatible with rotation of an RBG-bouquet. To see this, observe that u and v have rotation u, v, v, u at s. So any columns between u and v in both halves of the rotation must be in B. In particular, if z and z' are such loops, they must have rotation z, z', z', z. We conclude that there at most 3 loops in the cylinder.

Now consider all the loops not contained in the cylinder. Due to the rotation of s, they must all be in B. Suppose we cut the cylinder off the surface and fill x and y with discs. The new surface has genus $g - 2$. This surgery cannot have resulted in any additional contractible loops because the disc bounded by a contractible loop must contain x and y. This is prohibited by the rotation.

If two loops are newly homotopic, then the cylinder they bound contains x or y. Due to the rotation, if the cylinder contains one of x or y, it must also contain the other. Clearly there is only one such pair.

So, after removing the possible homotopic loop, we have a bouquet of non-contractible, pairwise non-homotopic loops, with each loop in B. By Corollary 3.2.2, there are at most $3(g-2) - 2$ loops. Including the loop made homotopic by removing the cylinder bounded by x and y, there are at most $3g - 7$ loops in the region bounded by x and y.

In total, there are at most 3 loops in the cylinder bounded by x and y, at most $3g-7$ loops in the other region, x and y on the boundary, r_1 which was originally cut, and r_2 which we contracted. We conclude that $\ell \leq (3g-7)+3+2+2 \leq 3g$.

Case 2 *Suppose g_1 is the contracted loop.* The argument is similar to the one used in Case 1. The only small difference is in Subcase (i) when the rotation of g_1, x, and y in the first half of the rotation is x, y, g_1 and in the second half is y, x, g_1. The discs bounded by g_1x and g_1y must still separate the surface in the same way and the result follows.

Case 3 *Suppose b_1 is the contracted loop.* If a loop u is newly contractible, then before contraction of b_1, the cycle formed by b_1u must have bounded a disc. Due to the rotation of s, this disc must contain either both of r_1 and r_1' or neither. If it contains neither, then u was homotopic to b_1 in Σ. If the disc contains both, then u must be the only such loop.

If there are no newly homotopic loops, after cutting and capping along r_1, contracting b_1, and removing the newly contractible loop u, we are left with $\ell-3$ non-contractible, non-homotopic loops on a surface of genus $g-1$. By induction, $\ell - 3 \leq 3g$ and we are done.

If x and y are newly homotopic, then we can proceed by an argument similar to Subcase (ii) above. There is a slight difference in the argument when analyzing the number of loops in the cylinder bounded by x and y. We see that a loop with both ends in the same side of the cylinder must be contractible and contain both r_1 and r_1', so there can only be one such loop. Otherwise, the argument is similar and that there at most 3 loops in the cylinder.

Considering the 3 loops in cylinder bounded by x and y, two on the boundary, and at most $3g-7$ loops in the other region, r_1 and b_1, we see that $\ell \leq 3g$. □

Chapter 4

Previous Results

In this section, we review the previous work done on crossing sequences. This includes:

(a) Širáň's original paper, in which he introduced crossing sequences and proved that every convex sequence is a crossing sequence;

(b) the paper of Archdeacon et al., which disproves Širáň's conjecture that every crossing is convex by showing that any sequence $a, b, 0$ is realized as a non-orientable crossing sequence; and

(c) the paper of DeVos et al., showing that every sequence $a, b, 0$ is realized as an orientable crossing sequence.

4.1 Širáň's Attempt at Characterization

In 1983, Širáň introduced the crossing function of a graph, the function that would later be known as the crossing sequence. It was Širáň's goal to characterize the set of sequences that could be the crossing sequences of some graph. He noted, as we have stated above, that any crossing sequence must be a finite sequence of non-negative integers that strictly decreases to zero. After looking at several small examples, he noted that all the graphs he had studied had convex crossing sequences.

Definition 4.1.1. A sequence, $a_0, a_1, a_2, \dots$, is *convex* if, for every $i \in \mathbb{N}$,

$$a_{i-1} - a_i \geq a_i - a_{i+1}$$

The main result of his paper is that convexity is a sufficient condition for a sequence being the crossing sequence of some graph.

Theorem 4.1.2. *Let $a_0, a_1, \ldots, a_n$ be a convex sequence of integers such that:*

(i) $a_n = 0$

(ii) for $1 \le i \le n$, $a_{i-1} > a_i$.

Then there exists a graph G with $a_0, a_1, \ldots, a_n$ as its crossing sequence.

The proof of this theorem uses a key result of graph embeddings which was proved by Battle et al. in 1962 [2]. Recall that the *genus of a graph*, denoted $\gamma(G)$, is the smallest integer g such that the graph can be embedded in a surface of genus g. A *block* of G is either a cut-edge or a maximal two-vertex-connected subgraph of G . The theorem given by Battle et al. states that the genus of a graph is additive over the genus of its blocks.

Theorem 4.1.3. *Suppose G is a connected graph containing k blocks $B_1, \ldots, B_k$ such that $E(G)$ is partitioned in to $E(B_1), \ldots, E(B_k)$. Then, $\gamma(G) = \sum_{i=1}^{k} \gamma(B_i)$.*

We now turn to the proof of Širáň's theorem.

Proof. Let $a_0, a_1, \ldots, a_n$ be a convex sequence with the properties (i) and (ii). Define a second sequence $b_1, b_2, \ldots, b_n$, where $b_i := a_{i-1} - a_i$. By property (ii) and convexity, $b_1 \ge b_2 \ge \ldots \ge b_n > 0$. We now construct the graph G, and will show that the weighted crossing sequence of G is $a_0, a_1, \ldots, a_n$.

We begin with n copies of $K_{3,3}$. For $1 \le i \le n$, let the i^{th} copy of $K_{3,3}$ have one edge of weight 1 and let the remaining edges have weight b_i. For each copy of $K_{3,3}$, pick one vertex that is incident to an edge of weight one and identify these vertices into one vertex, which we will call v. This produces a connected graph G, with a single cut-vertex v, and each of the original copies of $K_{3,3}$'s is a block of G. Since the genus of $K_{3,3}$ is one, Theorem 4.1.3 implies that G has genus n.

Now, for any $0 \le k \le n-1$, we wish to show that $\text{wcr}_k(G) \ge a_k$. We define $G(b_1, \ldots, b_j)$ to be the subgraph of G that contains the j copies of $K_{3,3}$ with edge weights $b_1, b_2, \ldots, b_j$. Since this subgraph has genus j, when $j \le k$, it can be embedded in the surface of genus k. When $j > k$, it is clear that there must be at least $j - k$ crossings. We will show that the weighted crossing number of this subgraph is $b_{k+1} + \cdots + b_j$.

First, suppose $j = k+1$. Then by Theorem 4.1.3, any drawing of $G(b_1, \ldots, b_j)$ on a surface of genus k must have a least one crossing. Moreover, two edges of weight 1 could not cross each other in any optimal drawing of G, since they would both be incident to vertex v and the drawing would not be optimal. So one of the other edges must be crossed. Therefore, the weighted crossing number is at least b_{k+1}.

Now, suppose that $k+1 < j \leq n$. Our induction hypothesis is that any drawing of the subgraph $G(b_1, \ldots, b_{j-1})$ contains $(j-1)-k$ crossings and the weight of these crossings is at least $b_{k+1} + \cdots + b_{j-1}$.

Choose an optimal drawing $\mathcal{D}$ of $G(b_1, \ldots b_j)$. By the induction hypothesis, $\mathcal{D}$ restricted to $G(b_1, \ldots, b_{j-1})$ contains at $j-1-k$ crossings with total weight $b_{k+1} + \cdots + b_{j-1}$. But $G(b_1, \ldots, b_j)$ has genus j and therefore must contain at least $j-k$ crossings. Thus, there is at least one crossing additional crossing that is not given in the induction hypothesis. As above, this cannot be a crossing of two edges of weight one, and so it must have weight at least b_j. So the weighted crossing number of $G(b_1, \ldots, b_j)$ is at least $b_{k+1} + \cdots + b_j$.

When $j = n$, $G = G(b_1, \ldots, b_n)$ and the weighted crossing number of G is at least $b_{k+1} + \cdots + b_n$. The reverse inequality is clear, as each block can be drawn with one crossing on the plane and each handle can save one weighted crossing. So we see that $\text{wcr}_k(G) = b_{k+1} + \cdots + b_n = a_k$. □

Širáň was unable to show that convexity was also a necessary condition for a sequence to be a crossing sequence. However, he was unable to find a counterexample and observations on "small graphs" left him to conclude the paper with the following conjecture:

Conjecture 4.1.4. *The crossing sequence of an arbitrary graph is convex.*

This conjecture has been the motivation to almost all recent results in the area of crossing sequences.

4.2 A Non-Convex Crossing Sequence

The conjecture remained open for over 15 years, until a counterexample was finally found by Archdeacon, Bonnington and Širáň [1]. They were able to provide examples of graphs with non-convex crossing sequences in both the orientable and non-orientable case. Their paper was also the first to address the relative non-convexity of a sequence. In particular, since a sequence is non-convex it must contain a non-convex jump, that is some i such that $a_{i-1} - a_i < a_i - a_{i+1}$. When this occurs, the ratio

$$\frac{a_{i-1} - a_i}{a_i - a_{i+1}}$$

gives some measure of the non-convexity of the jump. Presumably for ease of analysis, Archdeacon et al. worked with graphs that could be embedded in the double torus (or alternatively the Klein Bottle). Therefore, the goal was to show that for any $\varepsilon > 0$, there exists a graph G such that

$$\frac{cr_0(G) - cr_1(G)}{cr_1(G) - cr_2(G)} < \varepsilon$$

In the orientable case, this full generality was not reached. Their result proves the existence of a graph in which the number of crossings saved by adding the second handle approaches five times the amount saved by the first handle.

Their result for the case of non-orientable crossing sequences is more impressive. Although still limited to sequences of length three, they were able to show the following:

Theorem 4.2.1. *For any positive integers $a > b > 0$, there exists a graph with non-orientable crossing sequence $a, b, 0$.*

We will first describe their graph and then prove that it has the desired properties. We begin by fixing an arbitrary integer $n \geq 2$ and $q \geq 2n + 5$. The graph is the Cartesian product $C_{2n+2} \square P_q$ along with $n + 1$ additional edges. Let the path P_q contain vertices $\{1, \ldots, q\}$ and the cycle C_{2n+2} contain the vertices ℓ, then $t_1, \ldots, t_n$, followed by r and lastly $b_n, \ldots, b_1$. One way to picture the cycle is as a rectangle with the vertices labelled with an ℓ, t, r, and b on the left, top, right and bottom sides of the rectangle respectively.

We add n additional edges $\{(t_i, 1), (b_i, 1)\}$, for each $1 \leq i \leq n$. The last edge is $\{(\ell, 1), (r, 1)\}$.

Each of the edges in the Cartesian product of $C_{2n+2} \square P_q$ is thick and hence is uncrossed. The edge from $(t_1, 1)$ to $(b_1, 1)$ has weight a_0 and the remaining n edges have weight 1.

Lemma 4.2.2. $\widetilde{\mathrm{cr}}_0(G) = a_0 + (n - 1)$

Proof. Consider a drawing of G on the plane in which no thick edge is crossed. As we see in Figure 4.1, such a drawing must exist. Moreover, this picture shows that $\widetilde{\mathrm{cr}}_0(G) \leq a_0 + (n - 1)$

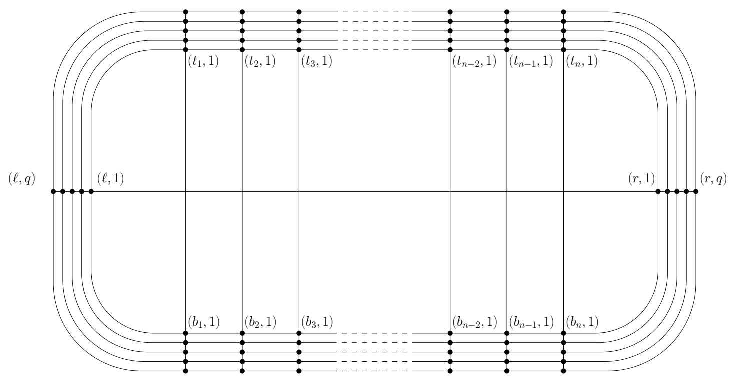

Figure 4.1: The graph described by Archdeacon, Bonnington, and Širáň.

Consider an optimal drawing of G. The cycles $C_{2n+2}\square 1$ and $C_{2n+2}\square 2$ separate the surface into three faces, each a disc. F_1 is bounded only by $C_{2n+2}\square 1$, F_2 is bounded only by $C_{2n+2}\square 2$, and F_3 is bounded by both curves.

Consider the edges $\{(\ell,1),(\ell,2)\}$, $\{(t_1,1),(t_1,2)\}$, $\{(r,1),(r,2)\}$, and $\{(b_1,1),(b_1,2)\}$. Each is uncrossed and these must be contained in F_3.

If $\{(\ell,1),(r,1)\}$ is contained in F_3, it must cross $\{(t_1,1),(t_1,2)\}$, $\{(b_1,1),(b_1,2)\}$, $C_{2n+2}\square 2$, or $C_{n+2}\square 1$. But this would result in $\{(\ell,1),(r,1)\}$ crossing a thick edge, so in any optimal drawing $\{(\ell,1),(r,1)\}$ is in F_1. Similarly, for any $j \in \{1,\dots n\}$, if $\{(t_j,1),(b_j,1)\}$ is in F_3, it must cross $\{(\ell,1),(\ell,2)\}$ or $\{(r,1),(r,2)\}$. So each such edge must also be in F_1, as otherwise a thick edge would be crossed.

For any j, if $\{(\ell,1),(r,1)\}$ is drawn on a disc with $\{(t_j,1),(b_j,1)\}$, then the two edges must cross. So the crossing number is at least the sum of the weights of the edges, $a_0 + (n-1)$. □

Lemma 4.2.3. $\widetilde{\mathrm{cr}}_1(G) = n-1$

Proof. Figure 4.1 shows G drawn on the plane. If a crosscap is used to locally remove the crossing between the edges $\{(\ell,1),(r,1)\}$ and $\{(t_1,1),(b_1,1)\}$, the crossing number is reduced to $n-1$. So $\widetilde{\mathrm{cr}}_1(G) \le n-1$.

Now assume that we have an optimal drawing of G in the projective plane. Consider the disjoint cycles $C_{2n+2}\square j$, for $j = 1,2,\dots,q$. Since disjoint non-contractible cycles in the projective plane must cross each other, at least $q-1$ of them must be contractible. Since $q \ge 2n+5$, there must be at least $2n+4$ uncrossed, non-contractible cycles labelled

$C_{2n+2} \square i_1, \ldots C_{2n+2} \square i_{2n+4}$ where $i_1 < \cdots < i_{2n+4}$. We wish to show that at least $n+2$ of these cycles are nested.

In order to prove this, we will need to consider an additional cycle. Let B denote the path $(t_1 \ell b_1 \ldots b_{n-1})$ in the cycle C_{2n+2}, while T is the disjoint path $(t_2 \ldots t_n r b_n)$. For any $j, k \in \{1, 2, \ldots q\}$, $T \square j$ and $T \square k$ can be connected on either end by subpaths of the paths $t_2 \square P_q$ and $b_n \square P_q$ to form a cycle $T_{j,k}$. Likewise, $B \square j$ and $B \square k$ are contained in a cycle $B_{j,k}$. $T_{j,k}$ and $B_{j,k}$ are disjoint, so at least one is contractible.

Let $j, k \in \{1, 2, \ldots 2n+4\}$ and let D_j and D_k be the discs bounded by the contractible cycles $C_{2n+2} \square i_j$ and $C_{2n+2} \square i_k$, respectively. If no two of the cycles D_j are disjoint, then we have all the cycles $C_{2n+2,i} \square i_j$ are nested, as required. Therefore, we may assume j and k are such that D_j and D_k are disjoint.

Since T_{i_j,i_k} and B_{i_j,i_k} are disjoint cycles and thick, one must be contractible. We assume that T_{i_j,i_k} is contractible, although the argument is symmetric for B_{i_j,i_k}. Let D_T be the disc bounded by T_{i_j,i_k}.

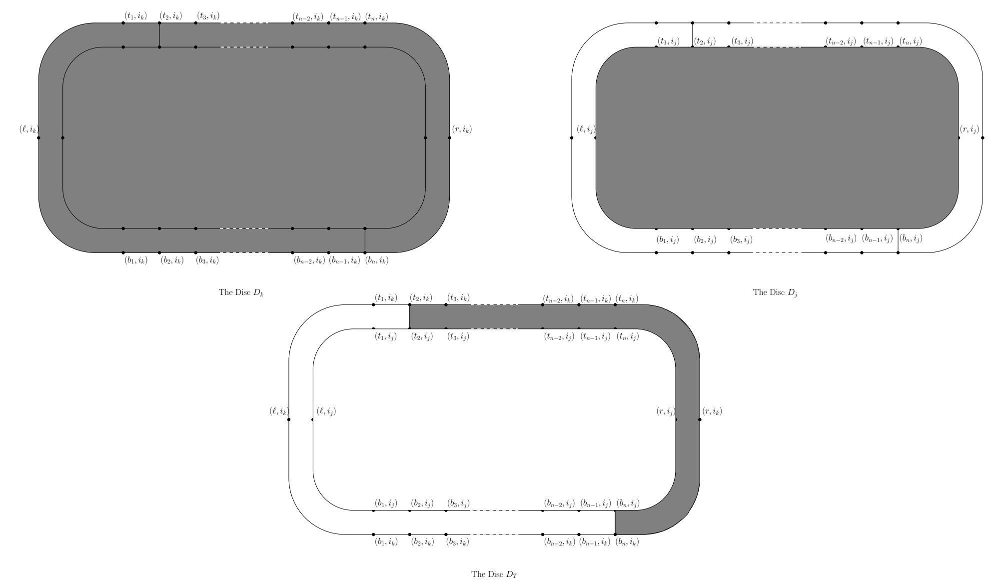

Figure 4.2: The discs D_k, D_j, and D_T. Here D_j and D_k are nested.

Case 1 *Suppose* $D_j \subseteq D_T$ *but* $D_k \nsubseteq D_T$. This is equivalent to saying that the path $B \square i_j$ is found in D_T but the path $B \square i_k$ is not. The points (r, i_j) and (r, i_k) do not share a face. In this case, the path $r \square P_q$ must cross a thick edge and the drawing is not optimal.

Case 2 *Suppose* $D_k \subseteq D_T$ *but* $D_j \nsubseteq D_T$. As in Case 1, the points (r, i_j) and (r, i_k) do not share a face, so the path $r \,\square\, P_q$ must cross a thick cycle. So, the drawing is not optimal.

Case 3 *Suppose* $(D_j \cup D_k) \subseteq D_T$ *or* $(D_j \cup D_k) \cap D_T = \emptyset$. Then, the vertices $(\ell, i_j), (r, i_j)$, (ℓ, i_k) and (r, i_k) do not share a common face.

Consider some cycle $C_{2n+2} \,\square\, i_m$ with $j < m < k$. For each of these 4 vertices, there exists a path, with interior disjoint from $(T_{i_j,i_k}) \cup (C_{2n+2} \,\square\, i_j) \cup (C_{2n+2} \,\square\, i_k)$, from the vertex to the cycle $C_{2n+2} \,\square\, i_m$. If such an m exists, then this path must cross a thick edge and the drawing is not optimal. Therefore j and k must be consecutive.

After embedding the cycles $(T_{i_j,i_k}) \cup (C_{2n+2} \,\square\, i_j) \cup (C_{2n+2} \,\square\, i_k)$, the vertices (ℓ, i_j) and (r, i_j) share exactly one face, D_j. We conclude that the cycle $C_{2n+2} \,\square\, i_m$ is contained in D_j for all $i_m < i_j$.

Similarly, the vertices (ℓ, i_k) and (r, i_k) share exactly one face, D_k. So, the cycle $C_{2n+2} \,\square\, i_{m'}$ is contained in D_k for all $i_{m'} > i_k$. Since there are $2n+4$ contractible cycles either $j \geq n+1$ or $k \leq n+2$. We will assume $j \geq n+2$, though the proof is symmetric for the other choice.

Consider some $m < j$; the disc D_m is contained in D_j. There are paths from (t_1, i_m) to (t_1, i_j) and from (b_1, i_m) to (b_1, i_j). These paths must separate the region $D_j \setminus D_m$ and, in particular, the vertices (ℓ, i_m) and (r, i_m) share only one face, D_m.

Now, consider the subpath of $\ell \,\square\, P_q$ between (ℓ, i_m) and (ℓ, i_{m-1}). It is internally disjoint from all the edges described in the previous paragraph and does not cross them. The same is true for the path between (r, i_m) and (r, i_{m-1}), so if (ℓ, i_{m-1}) and (r, i_{m-1}) lie on the same face, this face must be D_m.

But this is clearly true since (ℓ, i_{m-1}) and (r, i_{m-1}) are both contained in $C_{2n+2} \,\square\, i_{m-1}$, a thick cycle disjoint from the cycles previously embedded. So, the cycle $C_{2n+2} \,\square\, i_{m-1}$ is contained in D_m, for any $1 \leq m \leq j$. So the cycles $C_{2n+2} \,\square\, i_m$ are nested for $1 \leq m \leq j$.

Since the cycles are nested, we can conclude that the path $\ell \,\square\, P_q$ must cross each $C_{2n+2} \,\square\, i_j$ at the point (ℓ, i_j).

For each of the $n+2$ nested contractible cycles $C_{2n+2} \,\square\, i_0, \ldots, C_{2n+2} \,\square\, i_{n+1}$, we define a new cycle L_j. This cycle consists of the unique path of $C_{2n+2} \,\square\, i_j$ from (t_j, i_j) to (b_j, i_j) containing (ℓ, i_j) and the unique path from (b_j, i_j) to (t_j, i_j) that contains no edge in $C_{(2n+1)} \,\square\, i_k$ for any k. Note that the set of cycles $\{L_j\}_{j=1}^{n}$ is pairwise disjoint.

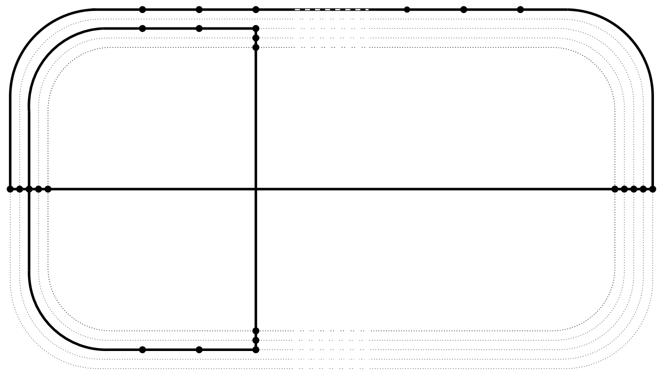

Figure 4.3: The cycles L_3 and T.

Define one additional cycle T consisting of the path in $C_{2n+2} \square i_{n+1}$ from (ℓ, i_{n+1}) to (r, i_{n+1}) containing (t_1, i_{n+1}) together with the paths $\ell \square P_q$, $\{(\ell, 1), (r, 1)\}$, and $r \square P_q$. Figure 4.3 shows a drawing of these cycles.

Suppose that L_i is contractible. Then, L_i intersects the cycle T at exactly one point and because the D_{i_j} are nested the cycles are not tangent to each other. Therefore, L_j and T cross at a vertex intersection. Since L_i is contractible, it must cross T somewhere else as well.

Alternatively, if L_i is non-contractible, it need not cross T. However, if L_j is also non-contractible, then L_i and L_j are disjoint non-contractible cycles in the projective plane and must cross.

We conclude that at least $n-1$ of the L_j's are crossed, and the crossing number is at least $n-1$. □

Using these two lemmas, we will now prove Theorem 4.2.1.

Proof of Theorem 4.2.1. Figure 4.4 provides an embedding of G on the Klein bottle, so $\widetilde{\text{cr}}_2(G) = 0$. From Lemma 4.2.2, we have that $\widetilde{\text{cr}}_0(G) = a_0 + (n-1)$ and Lemma 4.2.3 gives us that $\widetilde{\text{cr}}_1(G) = n-1$.

So the crossing sequence of G is $a_0 + (n-1), n-1, 0$. Given any positive integers $a > b > 0$, we can define $n := b+1$ and $a_0 := a - b$. Then the crossing sequence of G is $a, b, 0$. □

In particular, the sequence $a, b, 0$ can be arbitrarily non-convex. If $b = a - 1$, then we see that

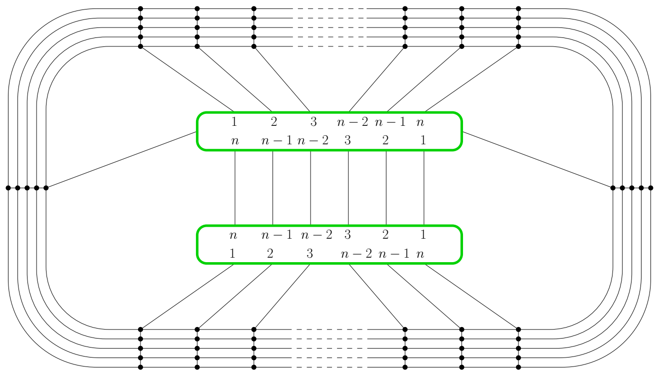

Figure 4.4: The graph G embedded in the Klein Bottle.

$$\frac{\widetilde{\mathrm{cr}}_0(G) - \widetilde{\mathrm{cr}}_1(G)}{\widetilde{\mathrm{cr}}_1(G) - \widetilde{\mathrm{cr}}_2(G)} = \frac{a - (a-1)}{(a-1) - 0} = \frac{1}{a-1}$$

Obviously, for any ε, there is a choice of a such that $\frac{1}{a-1} < \varepsilon$, meaning that this graph can contain an arbitrarily large non-convex jump. Although they were unable to prove it, Archdeacon et al. believed that this result was the same for the orientable case. In fact, their conjecture on the characterization of crossing sequences disregards convexity all together.

Conjecture 4.2.4. *Any sequence of integers that strictly decrease to zero, is the crossing sequence of some graph. It is also the non-orientable crossing sequence of a (different) graph.*

The article by DeVos et al. and this thesis are attempts to better understand this conjecture.

4.3 Arbitrary Crossing Sequences of Length Three

In 2011, DeVos, Mohar, and Šámal were able to prove the result that had eluded Archdeacon et al [3]. Theorem 4.3.1 provides the main result of their paper.

Theorem 4.3.1. *For any positive integers $a > b$, there exists a graph H with crossing sequence $a, b, 0$.*

We describe the graph H_n where $n \geq 3$. Let T and B be cycles $(t_0, \ldots, t_{n+1})$ and $(b_0, \ldots, b_{n+1})$, respectively, of length $n+2$. The vertices t_0 and b_0 are joined by a path of length 3 and the vertices t_{n+1} and b_{n+1} are as well. These paths are called the *left path* L containing points ℓ_1 and ℓ_2 and the *right path* R containing points r_1 and r_2, respectively. All of these edges are thick.

We will often refer to the thick edges as the *frame*. In any optimal drawing on a surface, it is clear that the frame must be embedded, as otherwise the crossing number is infinite. The largest cycle in the frame is of length $2n+8$ and contains every edge except $t_{n+1}t_0$ and $b_{n+1}b_0$. We will denote this cycle by S and refer to it often.

For each $1 \leq i \leq n$, there is a *column* c_i which is a path from t_i to b_i of length 2. Let m_i be the midpoint of column c_i. A column c_i is *even* if i is even and is *odd* if i is odd. Finally, the *odd row* R_o is the path from ℓ_1 to r_1 containing (in the natural order) the midpoint of each odd column. The *even row* R_e is the path from ℓ_2 to r_2 containing (in the natural order) the midpoint of each even column.

We specify that the vertices located at the intersection of a row and column are rigid vertices. The rotation is prescribed so that the row and column cross at the given vertex. In particular, the row is not merely tangent to the column. Figure 4.5 provides a drawing of the graph H_n.

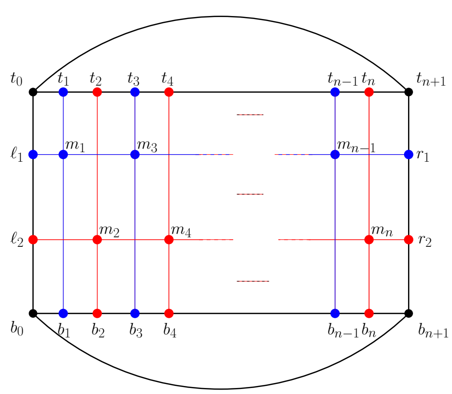

Figure 4.5: The graph H_n introduced by DeVos, Mohar, and Šámal.

For our discussion of the columns of H_n, we will often need to discuss their homotopy. In order to do so, we introduce a slightly modified definition of homotopy.

Two columns c_i and c_j, with $i < j$, are *homotopic* if the cycle contained in $\{t_i \ldots t_j\} \cup \{b_i \ldots b_j\} \cup c_i \cup c_j$ bounds a closed disc. A column c_i is *contractible* if the cycle contained in $\{t_0 \dots t_i\} \cup \{b_0 \dots b_i\} \cup L \cup c_i$ or the cycle contained in $\{t_i \dots t_{n+1}\} \cup \{b_i \dots b_{n+1}\} \cup R \cup c_i$ bounds a closed disc.

This definition is similar to the standard definition of homotopy and, in particular, any two homotopic columns are also homotopic in the traditional sense. However, the converse is not necessarily true. For example, two contractible columns need not be homotopic, as there are up to two distinct homotopy classes for contractible columns.

The following lemma establishes that the homotopy types of the columns will often be enough to guarantee crossings. The lemma supposes an optimal drawing of H_n. In order for this to be meaningful, we note that Figure 4.5 shows a drawing in H_n with n crossings. So $\mathrm{cr}_g(H_n) \leq n$.

Lemma 4.3.2. *Suppose H_n is drawn optimally on a surface Σ. If there are $j \geq 2$ consecutive columns, $c_i, c_{i+1} \dots, c_{i+j-1}$ that are pairwise homotopic, then these columns contain at least j crossings.*

Proof. Suppose $j = 2$. Since c_i and c_{i+1} are homotopic, they bound a disc. Since there is path from L to R disjoint from the cycle bounded by c_i and c_{i+1} and consisting of thick edges, either L and R are both inside this disc or both outside.

Each row has a vertex in the boundary of the this disc. The rotation of this vertex is prescribed in such a way that the row must cross the boundary of the disc at the vertex. Therefore, it must cross the boundary of the disc at another point. Thus, both rows cross the boundary of the disc, as required.

Suppose $j = 3$ and that c_i, c_{i+1}, and c_{i+2} are consecutive homotopic columns. As above, each row must cross the boundary of the disc bounded by c_i and c_{i+1}.

If c_{i+1} is not contained in the disc bounded by c_i and c_{i+2}, then it must cross the boundary of this disc twice, yielding two additional crossings. Otherwise, c_{i+1} is contained in the disc bounded by c_i and c_{i+3}. In this case, the row that contains the vertex in c_2 must cross this disc twice. This yields at least 3 crossings in total, as needed.

Suppose now that $j \geq 4$ and let $c_i, \dots, c_{i+j-1}$ be consecutive homotopic columns. By induction, $c_i, \dots, c_{j-3}$ contains $j - 2$ crossings. Additionally, each row crosses the disc bounded by c_{i+j-2} and c_{i+j-1}, completing the proof. □

With this lemma, we can now begin the proof of Theorem 4.3.1 by determining the crossing numbers on the sphere and the double torus.

Lemma 4.3.3. $\mathrm{cr}_0(H_n) = n$.

Proof. We begin by embedding S on the sphere. This separates the sphere into two discs, D_1 and D_2. If the edges $t_{n+1}t_0$ and $b_{n+1}b_0$ are not contained in the same disc, then each column must cross the frame and the crossing number is infinite.

We will assume that they are both in D_2. All the columns, and therefore both rows, must be in D_1. Each column separates D_1 into two regions, with L and R on the boundary of opposite regions. Thus, each even column must be crossed by the odd row and each odd column must be crossed by the odd row. The crossing number is at least n. Figure 4.5 shows that equality must hold. □

Lemma 4.3.4. $\mathrm{cr}_2(H_n) = 0$.

Proof. Begin by drawing S on the sphere. As in Lemma 4.3.3, this separates the sphere into two discs, D_1 and D_2. We will draw all the odd columns and R_o in D_1 without crossings, and do the same for the even columns and R_e in D_2. We can then use the new handles to draw each of $t_{n+1}t_0$ and $b_{n+1}b_0$, leaving us with an embedding on the double torus. □

Unfortunately, determining the crossing number of H_n on the torus is much more complicated and relies heavily on analyzing how the thick edges are embedded.

Lemma 4.3.5. *If* $n = 3$ *or* $n \geq 5$, *then* $\mathrm{cr}_1(H_n) = n - 1$.

The proof offered by DeVos et al. first specifies the ways in which the frame can be embedding in an optimal drawing and proceeds to attack each case individually. We will follow their method of proof and look at possible embeddings of the thick edges of H_n. To do so, we first introduce several properties of the embedding.

Claim 4.3.6. *Suppose the cycle T (or B) is contractible in an optimal drawing of H_n on the torus and bounds a disc D. Then, for each i, the column c_i is disjoint from D.*

Proof. Suppose T is contractible and bounds a disc D. If the cycle B is in D, then each column and therefore each row is also in D. Since each row is in D, the left and right paths much also be in D, so we have a planar drawing of H_n and there at least n crossings.

We now suppose B is not in D. The column c_i has exactly one endpoint in T. Since T is thick and uncrossed, if c_i was not disjoint from D it would cross the boundary of the disc. This results in an infinite crossing number and the drawing is not optimal, so the column must be disjoint from D.

In particular, each column must be incident to the cycle T on the side opposite from the ends of the edge $t_{n+1}t_0$. The proof is similar if B is contractible. □

Claim 4.3.7. *Suppose the cycle T (or B) is homotopic to S in an optimal drawing of H_n on the torus. Then, together S and T (or B) bound a disc D and each row and column is disjoint from D.*

Proof. Suppose T and S are homotopic. So $t_{n+1}t_0$ and $L \cup (B \setminus b_{n+1}b_0) \cup R$ form a cycle that bounds a disc D.

If the path $t_0t_1 \dots t_{n+1}$ is contained in D, then each columns and each row is contained in D. The cycle B can be drawn as a contractible cycle to obtain a planar drawing of H_n with no additional crossings. Such a drawing has at least n crossings, so we can assume that the path $t_0t_1 \dots t_{n+1}$ is not contained in D.

For each i, c_i is incident to cycle exactly once. If the column c_i were not disjoint from D, it would cross the boundary of the disc, yielding an infinite crossing number.

Since each column is disjoint from D, if R_e is not disjoint from D it must cross the boundary. The same is true for R_o. We conclude that each row and column is disjoint from D.

In particular, each column is incident to B on the boundary of D. Since the column is disjoint from D, we know that the column must be incident to B on the opposite side of S as $t_{n+1}y_0$. □

Using the above claims, we can now classify the many possibilities for the drawing of H_n on the torus. In particular, we will classify how the frame can be embedded.

S is Contractible S separates the torus into two faces, D_1, a disc, and D_2. If the interior of H_n is completely drawn in D_1, then each column is crossed, resulting in n crossings. From Lemma 4.3.3, $\mathrm{cr}_1(H_n) \leq n-1$, so the interior cannot be contained in D_1.

If $t_{n+1}t_0$ and $b_{n+1}b_0$ are both in D_2, then T and B are disjoint cycles in D_2. If either is contractible, then the interior of H_n is in D_1, a contradiction, so we may assume both are non-contractible. Cutting and capping along T leaves a new surface Σ' homeomorphic to the sphere. In Σ', B is contractible. Thus, the interior of the H_n must have been drawn in D_1. This is a contradiction, so $t_{n+1}t_0$ or $b_{n+1}b_0$ must be drawn in D_1.

Suppose exactly one of $t_{n+1}t_0$ and $b_{n+1}b_0$ is drawn in D_1. We will suppose it is $b_{n+1}b_0$, but the argument for the other choice is symmetric. Since B is contractible, the entire interior of the graph must be drawn in D_2. In this case, $t_{n+1}t_0$ can be redrawn in D_1 with no additional crossings. So for any optimal drawing with the exactly one of $t_{n+1}t_0$ and $b_{n+1}b_0$ drawn in D_2, there is an optimal drawing with both drawn in D_1.

Therefore, we only need to consider drawing with the both those edges drawn in D_1. Drawings with the frame embedded in this way will be *Type* A.

S is Non-Contractible

Case 1 *Suppose the ends of* $t_{n+1}t_0$ *and* $b_{n+1}b_0$ *are all incident to the same side of* S*.* If both T and B are contractible, we will call the drawing *Type* B.

Since the ends of $b_{n+1}b_0$ are incident to the same side of S, if B is non-contractible it must be homotopic to S. By Claim 4.3.7, $S \cup B$ must separate the surface and one region of this separation is a disc. This disc contains both ends of T, so T must be contractible. When the frame is embedded in this way, we will say it is a *Type* C embedding.

Alternatively, if we have the symmetric case where T is not contractible, we call the embedding *Type* C$'$.

Case 2 *Suppose* $t_{n+1}t_0$ *has both ends incident to one side of* S *and* $b_{n+1}b_0$ *has both ends incident to the other.* Let $t_{n+1}t_0$ be incident to the side S_1 and $b_{n+1}b_0$ incident to the side S_2.

If T is contractible, each column must be incident to T on the side S_2, by Claim 4.3.6. If T is non-contractible, each column must be incident to B on the side S_2, by Claim 4.3.7. Similarly, if B is contractible, each column must be incident to B on the side S_1. If B is non-contractible, each column must be incident to T on the side S_1.

From this, we conclude that T and B are either both contractible or both non-contractible. If both are non-contractible, then one of the discs discussed in Claim 4.3.7 will not be disjoint from the rows and the drawing is not optimal. If both are contractible, the embedding will be called *Type* D.

Case 3 *Suppose* $b_{n+1}b_0$ *has ends incident to the same side of* S *and* $t_{n+1}t_0$ *has ends incident to opposite sides of* S*.* Let B be incident to the side S_1. If B is non-contractible, consider the disc described in Claim 4.3.7. The edge $t_{n+1}t_0$ must have one end in this and hence disc cross B. In this case the drawing is not optimal, so B must be contractible.

The embedding is *Type* E if $t_{n+1}t_0$ is incident to side S_1 at t_0 and is *Type* E$'$ if it is incident to S_1 at t_{n+1}.

Case 4 *Suppose* $t_{n+1}t_0$ *has ends incident to the same side of* S *and* $b_{n+1}b_0$ *has ends incident to opposite sides of* S*.* This is symmetric to Case 3, so T must be contractible. The embedding is *Type* E$''$ if $b_{n+1}b_0$ is incident to side S_1 at b_0 and is *Type* E$'''$ if it is incident to S_1 at b_{n+1}.

Case 5 *Suppose both $t_{n+1}t_0$ and $b_{n+1}b_0$ have ends incident to opposite sides of S.* In this case, T and B must both be non-contractible and, since they are uncrossed, homotopic. Since S is homotopic, ℓ_1 and r_1 lie in different cylinders and hence R_e must cross $T \cup B$.

Figure 4.6 shows the possible embeddings of thick edges in the torus, with outer rectangle representing the tours. For the drawings (B), (C), (C′), (D), (E), (E′), (E″), and (E‴), the horizontal (identified) edge of the rectangle is the non-contractible cycle S.

From the previous discussion, we see that if we can establish a minimum number of crossings for each of the 9 cases discussed above, then the crossing number will be at least that minimum. We now proceed to show that for every possibility, there are at least $n-1$ crossings.

Proposition 4.3.8. *If $\mathcal{D}$ is a Type* A *drawing of H_n on the torus, then there are at least $n-1$ crossings in $\mathcal{D}$.*

Proof. In a Type A drawing, S is contractible. Contract it to a point s. The columns form a bouquet with rotation as described in Theorem 3.2.1, with each loop in B.

If a column is contractible, it is crossed by a row. Therefore, in order for the drawing to contain less than n crossings, at least one column must be non-contractible. Suppose c_i is non-contractible.

Consider any other non-contractible column c_j. Due to the rotation of the bouquet, c_j is homotopic to a non-contractible cycle, disjoint from c_i. On the torus, disjoint non-contractible cycles are either homotopic or they cross [4, Prop. 4.2.6]. So c_i and c_j are either homotopic or cross.

Suppose there are p contractible columns, q non-contractible columns not homotopic to c_i, and r columns homotopic to c_i. Clearly, $p+q+r = n-1$, since they account for all columns except c_i.

Each contractible column is crossed by a row, so these p columns contribute at least p crossings. Each non-contractible column not homotopic to c_i crosses c_i, and hence these q columns contribute at least q crossings.

If $r = 0$, we are done, so suppose $r \geq 1$, consider the homotopic columns $c_{j_0}, \ldots, c_{j_r}$. If they are consecutive, then by Lemma 4.3.2, these $r+1$ columns contribute at least $r+1$ crossings.

If these columns are not consecutive then c_{j_0} and c_{j_r} bound a cylinder containing a column not homotopic to either. This column must cross the boundary of the cylinder

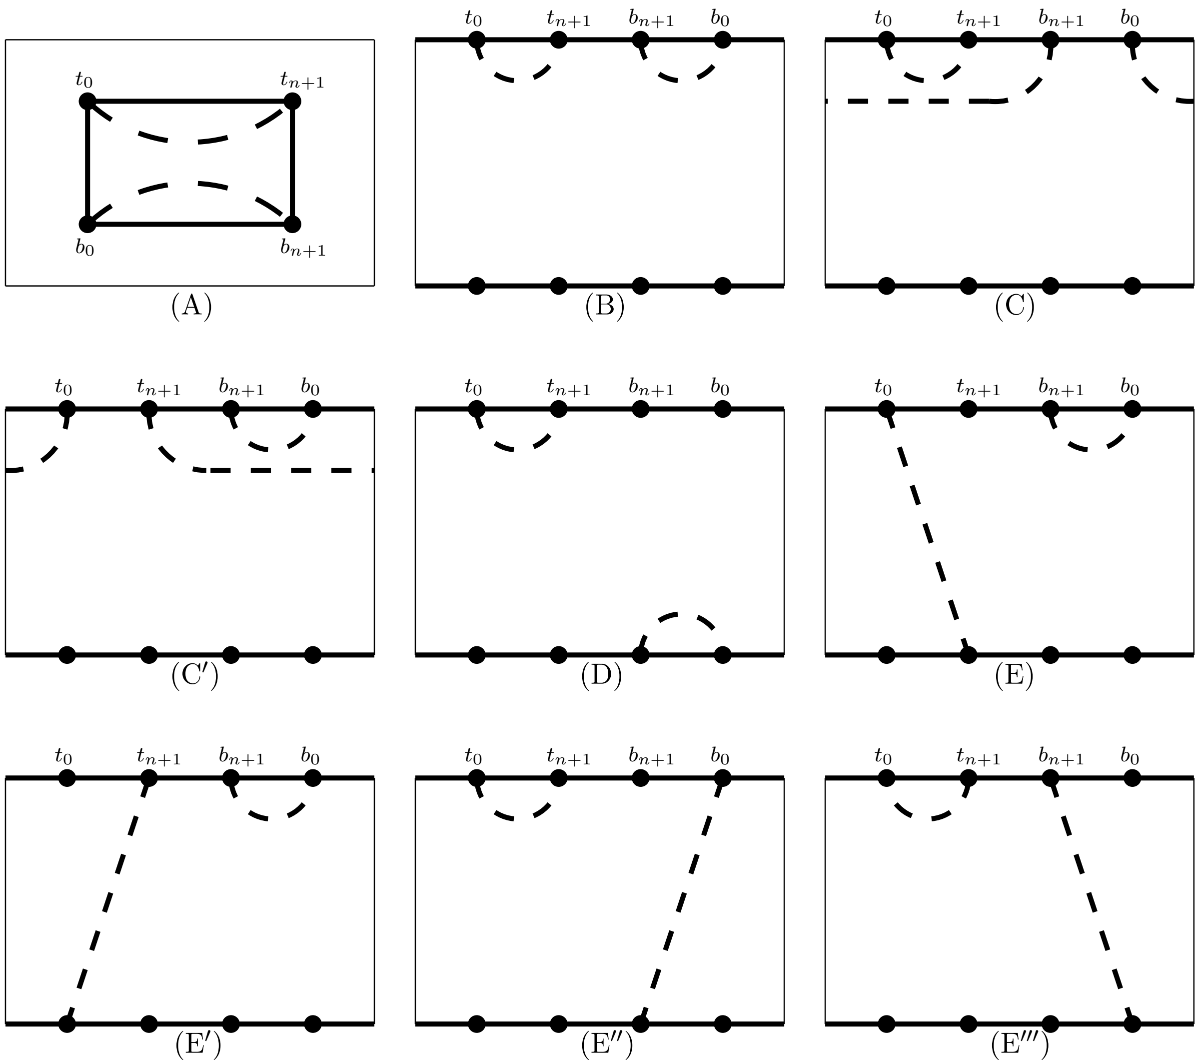

Figure 4.6: Possible embeddings of the thick edges of H_n on the torus.

twice, and one of these crossings is previously unaccounted for. So either c_{j_0} or c_{j_r} is crossed and we can remove it to get a drawing with one fewer crossing and one fewer column homotopic to c_i. By induction, the $r+1$ columns contain at least r crossings.

So, the r columns contribute at least r crossings. So the crossing number is at least $p+q+r=n-1$. □

Proposition 4.3.9. *If $\mathcal{D}$ is a Type* B *drawing of H_n on the torus with $n \neq 4$, then there* *are at least* $n-1$ *crossings in* $\mathcal{D}$.

Proof. Contract the edges of $\{t_0 \dots t_{n+1}\} \cup L \cup \{b_0 \dots b_{n+1}\}$ to a point s. The columns $c_1, \dots, c_n$ and the path R form a bouquet with $n+1$ loops.

Since drawing is Type B, we know that the path R is non-contractible and, since it is thick, uncrossed. Due to the rotation of s, if any column is non-contractible, it is homotopic to a cycle disjoint from R. Since disjoint non-contractible cycles on the torus either cross or are homotopic, any non-contractible column must be homotopic to R.

Adopting the notation of DeVos et al., we label a column ℓ if it is contractible and label it r if it is homotopic to R. We proceed by induction. The base cases are $n=3$ and $n=6$.

For $n=3$, we wish to show that there are at least 2 crossings. If c_2 is contractible and c_1 is non-contractible, c_1 and c_2 cross twice. If both c_1 and c_2 are contractible, then together they bound a disc. Each row has a vertex in the boundary of the this disc. The rotation of this vertex is prescribed in such a way that the row must cross the boundary of the disc at the vertex. Therefore, it must cross the boundary of the disc at another point. This is true for each row, yielding two crossings.

The argument is analogous if instead c_2 is homotopic to R; in this case we consider c_3 rather than c_1.

For $n=6$, let j be the smallest integer such that c_j is labelled r. If no such column exists, then all are contractible and there are at least 5 crossings. Every for every $k>j$, either c_k is marked r as well or c_k crosses c_j twice. If $j \geq 3$, then the columns $c_1, \dots c_{j-1}$ contain $j-1$ crossings.

Case 1 *Suppose* $j \leq 2$. Consider the columns $c_{j+1}, \dots, c_6$. At most 2 are labelled ℓ and each of these cross c_j twice. If there are exactly 2, then these must be c_{j+1} and c_{j+2}, as otherwise there would be more than 5 crossings. But in this case, c_{j+3} and c_{j+4} must be homotopic and contribute 2 crossings, leaving at least 5 in total.

If there are no columns marked ℓ, then $c_2, \dots, c_6$ are all homotopic and contain 5 crossings.

If only one of $c_{j+1}, \dots c_6$ is labelled ℓ, then 4 of c_2, c_3, c_4, c_5, c_6 are labelled r. If c_4 is labelled ℓ, then c_2 and c_3 are homotopic and c_5 and c_6 are homotopic. By Lemma 4.3.2, each pair contributes 2 crossings. When added to the two crossings from the column labelled ℓ, we see there are at least 5 crossings.

If c_4 is not labelled ℓ, then either c_2, c_3, c_4 or c_4, c_5, c_6 is a set of 3 consecutive columns all labelled r. By Lemma 4.3.2, one of these sets must contribute at least 3 crossings. The two additional crossings from the column labelled ℓ give at least 5 crossings.

Case 2 *Suppose* $j = 3$. The columns c_1 and c_2 must be homotopic and contribute 2 crossings. Since each column labelled ℓ crosses c_3 twice, there are at least 5 crossings if more than one of c_4, c_5 and c_6 are labelled ℓ.

If none of c_4, c_5, c_6 are labelled ℓ, they are homotopic and together contribute three crossings. If exactly one of c_4, c_5, c_6 is marked ℓ, then it is crossed twice and either c_4 is homotopic to c_3 or c_5 is homotopic to c_6. In either case there are 4 crossings from the final 4 columns.

Case 3 *Suppose* $j = 4$ *or* $j = 5$. c_1, c_2, and c_3 are all homotopic and contribute 3 crossings. Either c_{j+1} is labelled ℓ and crosses c_j twice or c_{j+1} is homotopic to c_j and each row must cross one of the two. In either case there are 5 crossings.

Case 4 *Suppose* $j = 6$. In this case, $c_1, \ldots, c_5$ are all contractible and contain 5 crossings.

Now suppose $n = 5$ or $n \geq 7$, we will attempt to find two consecutive columns that together contain two crossings. If such columns exist, we can remove them and obtain a drawing of H_{n-2} two fewer crossings. The result then follows by induction. Let j be the smallest number such that c_j is marked r.

If such a j doesn't exist, then all the columns are contractible and by Lemma 4.3.2 contain n crossings.

If $j = 1$, then either c_2 is labelled ℓ and crosses c_1 twice or c_2 is also marked r and is homotopic to c_1. In this case, Lemma 4.3.2 shows that c_1 and c_2 contain two crossings. We can remove c_1 and c_2 and proceed by induction.

If $j = 2$ and all remaining columns are labelled r, then $c_2, \ldots, c_n$ are homotopic and contain $n - 1$ crossings, and we are done. Otherwise, let c_i be be the first column after c_j to be marked ℓ. It crosses c_{i-1} twice, so we can remove both c_{i-1} and c_i and proceed by induction.

If $j \geq 3$ then two columns c_1 and c_2 must be labelled ℓ and therefore homotopic. Then c_1 and c_2 contain two crossings, so we can remove them and proceed by induction. □

Proposition 4.3.10. *If* $\mathcal{D}$ *is a Type* C *drawing of* H_n *on the torus, then there are at least* $n - 1$ *crossings in* $\mathcal{D}$.

Proof. In this type of drawing, both rows must be incident to the same side of S, namely the side not incident with $t_{n+1}t_0$ and $b_{n+1}b_0$. Since the drawing is minimal, the cycle contained in $R \cup L \cup R_e \cup R_o$ cannot bound a disc, as otherwise each column would cross the cycle.

So either the cycle contained in $\{t_0, \ldots, t_{n+1}\} \cup R \cup L \cup R_e$ bounds a disc or the the cycle contain in $\{t_0, \ldots, t_{n+1}\} \cup R \cup L \cup R_o$ bounds a disc. In the former case, each odd column must cross the disc and in the latter, each even column must cross the disc. We will assume the latter case and look for additional crossings containing the odd columns, although the other case is similar.

Since the rows are incident to the same side of S, any column that is contractible must be crossed by a row. In particular, all such odd columns are crossed.

Any remaining odd columns are incident to opposite sides of S. Suppose c_i and c_k are such odd columns, with $i < k$. If they are homotopic, then the two columns cross and, along with the path $b_i, \ldots, b_k$, separate the surface into three regions. If they are not homotopic, then the cycle formed by $\{b_i, \ldots b_k\} \cup c_k \cup \{t_k, \ldots, t_{n+1}, t_0, \ldots, t_i\} \cup c_i$ separates the surface.

In either case, there must exist an even column c_j between c_i and c_k. The column has endpoints to different faces of the separation described in the previous paragraph. The column c_j must cross the boundary of the region bounded by c_i and c_k. Hence c_j crosses at least one of c_i and c_k.

Moreover, if there are p odd columns incident to opposite ends of S, then there must be $p-1$ crossings between odd and even columns. We can conclude that there are at least $n-1$ crossings. □

Note that the proof for when H_n is a Type C′ drawing follows a symmetric argument.

Proposition 4.3.11. *If $\mathcal{D}$ is a Type* D *drawing of H_n on the torus, then there are at least $n-1$ crossings in $\mathcal{D}$.*

Proof. Cut along S to obtain a sphere with two holes. These holes are bounded by two copies of S, which we will call S_1 and S_2.

We can suppose $b_{n+1}b_0$ is incident to S_1 and $t_{n+1}t_0$ is incident to S_2. Every column must be incident to T at S_1 and incident to B at S_2

Assume that $t_0, \ldots, t_{n+1}$ are in clockwise order in S_1. Then, since the original surface was orientable, $b_0, \ldots, b_{n+1}$ must be also be ordered in clockwise direction in S_2.

This proof will be by induction. Suppose $n = 3$. It must be shown that there are at least 2 crossings.

Suppose c_1 and c_2 cross. Then the edge t_1t_2 on S_1 together with the portion of c_1 and c_2 above the crossing point must bound a disc. Similarly, the edge b_1b_2 together with the portions of c_1 and c_2 below the crossing point bound a disjoint disc. The row R_e crosses the boundary of one of these discs at a vertex and must cross it at a different point. The same is true for R_o. Hence there are two crossings.

Similarly, if c_2 and c_3 cross, there must also be at least two crossings. We can now suppose that c_2 crosses neither c_1 nor c_3.

If c_1 and c_3 cross, then there is a disc bounded by $t_1t_2t_3$ in S_1 and the portions of c_1 and c_3 up to the crossing point. This disc contains exactly one end of c_2 and, hence, c_2 must cross c_1 or c_3.

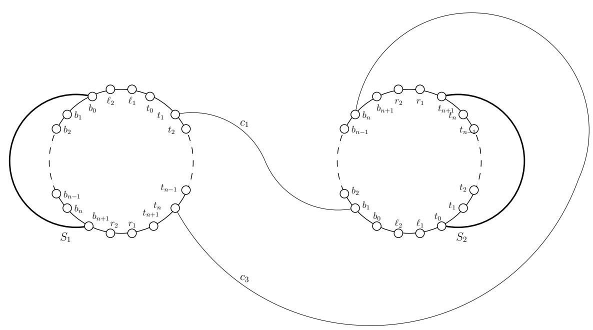

Figure 4.7: A drawing of the uncrossed columns c_1 and c_n in a Type D drawing of H_n, after cutting the torus along S.

Alternatively, if c_1 and c_3 do not cross, then the cycle $(\{t_1t_2t_3\} \cap S_1) \cup c_1 \cup \{(S_2 \setminus \{t_1t_2t_3\}) \cup c_3$ separates the surface. Each edge of this separation contains exactly one end of c_2. We conclude that c_2 must cross c_1 or c_3 and that there are at least 2 crossings.

Now, let $n \geq 4$. If c_1 or c_n is crossed, then we can delete it and the result holds by induction.

Suppose neither c_1 nor c_n is crossed. In particular, they do not cross each other. Due to the orientation of the S_1 and S_2, the cycle $(\{t_1t_2t_3\} \cap S_1) \cup c_1 \cup \{(S_2 \setminus \{t_1t_2t_3\}) \cup c_3$ separates the surface. However, each region contains exactly one end of the column c_2. So c_2 must cross either c_1 or c_n, a contradiction.

Therefore, either c_1 or c_n is crossed and the proof is complete by induction. □

Proposition 4.3.12. *If $\mathcal{D}$ is a Type* E *drawing of H_n on the torus, then there are at least $n-1$ crossings in $\mathcal{D}$.*

Proof. Again we proceed by induction. Note first that all columns are incident to B on the same side of S. As before a column may be contractible or non-contractible. However, because of the caps, any two contractible columns are homotopic. Likewise, any two non-contractible columns are homotopic. Additionally, because of the rotation, if two columns are non-contractible, they cross each other.

Let $n = 3$. By Lemma 4.3.2, if c_1 and c_2 are homotopic, there are at least two crossings. Similarly, if c_2 and c_3 are homotopic, there are also at least two crossings. So either c_2 is the only contractible loop or c_2 is the only non-contractible loop.

If c_2 is contractible, then c_1 must be non-contractible and it crosses the boundary of the disc bounded by c_2. Additionally, c_3 must also be homotopic and therefore cross c_1. So there at least two crossings.

So we may suppose c_2 is non-contractible. In this case c_3 is contractible and must cross c_2. Now, one end of R_e must be contained in the region bound by $t_{n+1}t_0$ and c_2. If R_e has both ends in the same side of S, it must cross c_3. If it has both ends in the opposite sides of S then, together with S and $t_{n+1}t_0$, it separates the surface. In particular, R_o must have one end in each region of this separation and contribute one crossing.

We conclude that when $n = 3$, there are at least 2 crossings.

Now, let $n \geq 4$. Suppose first that c_n is non-contractible. If it is crossed, we can remove it and obtain the result by induction. If it is uncrossed, then the remaining $n - 1$ columns must be contractible. By Lemma 4.3.2, these columns must contribute $n - 1$ crossings.

Alternatively, suppose c_n is contractible. Then, c_n together with S bounds a disc containing at least one end of $c_1, \ldots, c_{n-1}$. If one of the remaining loops is non-cotnractible, it must cross c_n. So we can remove c_n, have a drawing with one fewer column, one fewer crossing and be done by induction. If all the columns are contractible, by Lemma 4.3.2 they must contribute at least n crossings. □

The proofs for when the drawing of H_n is Type E$'$, E$''$, or E$'''$ follow a symmetric argument. Through Propositions 4.3.8, 4.3.9, 4.3.10, 4.3.11, and 4.3.12, we see that the crossing number of H_n on the torus is at least $n - 1$, for appropriate n. This is enough to prove Lemma 4.3.5.

Moreover, Lemma 4.3.3, Lemma 4.3.5, and Lemma 4.3.4 show that the crossing sequence of H_n is $n, n - 1, 0$ for $n \neq 4$.

To achieve the full generality of Theorem 4.3.1, observe that the proof of Lemma 4.3.3 is stronger than just showing the crossing number is n. It also shows that every column in H_n is crossed by a row.

Suppose we define a weighted graph $H_{n,k}$ to be equivalent to H_n except the column c_2 has weight k for $k \in \mathbb{N}$. All thick edges in H_n are thick in $H_{n,k}$ and the remaining columns and rows have weight 1.

The planar crossing number of $H_{n,k}$ is $k + n - 1$. The toroidal crossing number is still $n - 1$ and $H_{n,k}$ embeds in the double torus. For given integers $a > b \geq 2$ with $b \neq 3$, let $n := b + 1$ and $k := a - b$. The crossing sequence of $H_{n,k}$ is $a, b, 0$.

In particular, If $b = 1$, any crossing sequence $a, 1, 0$ is convex and, by Theorem 4.1.2, there exists a graph with crossing sequence $a, 1, 0$. Moreover, if $b = 3$, the sequence $a, 3, 0$ is convex for any $a \geq 6$. So all that remains to be shown is a graph with crossing sequence $5, 3, 0$ and a graph with crossing sequence $4, 3, 0$.

Devos et al. introduce a new graph H_3^+ for these two sequences. Begin with the graph H_3. Subdivide the edges $\{\ell_2, b_0\}, \{m_1, b_1\}, \{m_3, b_3\}$, and $\{r_2, b_4\}$. We add an additional row, R_3 containing the four new points obtained through these subdivisions.

For ease of notation, we now refer to the original odd row as R_1 and the original even row as R_2. The vertex at the intersection of row R_i and column c_j is $m_{i,j}$. The new vertices in columns c_1 and c_3 are rigid, with the rotation such that the row and column are not tangent. Figure 4.8 shows a drawing if H_3^+.

The proofs that $\text{cr}_0(H_3^+) = 4$ and $\text{cr}(H_3^+) = 0$ follow very closely to Lemma 4.3.3 and Lemma 4.3.4, respectively. The case for the torus makes use of Lemma 4.3.5 for H_3, but requires some additional insight.

Claim 4.3.13. *The crossing sequence of H_3^+ is $4, 3, 0$.*

Proof. Consider an optimal drawing of H_3^+ on the plane. The cycle S separates the surface into two discs and one of these discs contains the edge t_4t_0. The other disc must then contain c_1, c_2, and c_3 and, hence, R_1, R_2, and R_3.

Due the rotation of vertices in S, c_2 must cross R_1 and R_3, while R_2 must cross c_1 and c_3. So the crossing number is at least 4. Moreover, c_1 is crossed at least once. Figure 4.8 shows a drawing on the plane with 4 crossings, so equality must hold.

Note that the subgraph $H_3^+ \setminus (\{t_0, t_4\} \cup \{b_0 b_4\})$ is planar. Given a drawing of this subgraph on the sphere, two handles can be added to the surface so that $\{t_0 t_4\}$ and $\{b_0 b_4\}$ can be drawn without introducing any crossings. This gives and embedding of H_3^+ on the double torus.

Now, consider an optimal drawing of H_3^+ on the torus. If R_1 is crossed, remove it to obtain a drawing that is isomorphic to a subdivision of H_3. This subdivision contains at

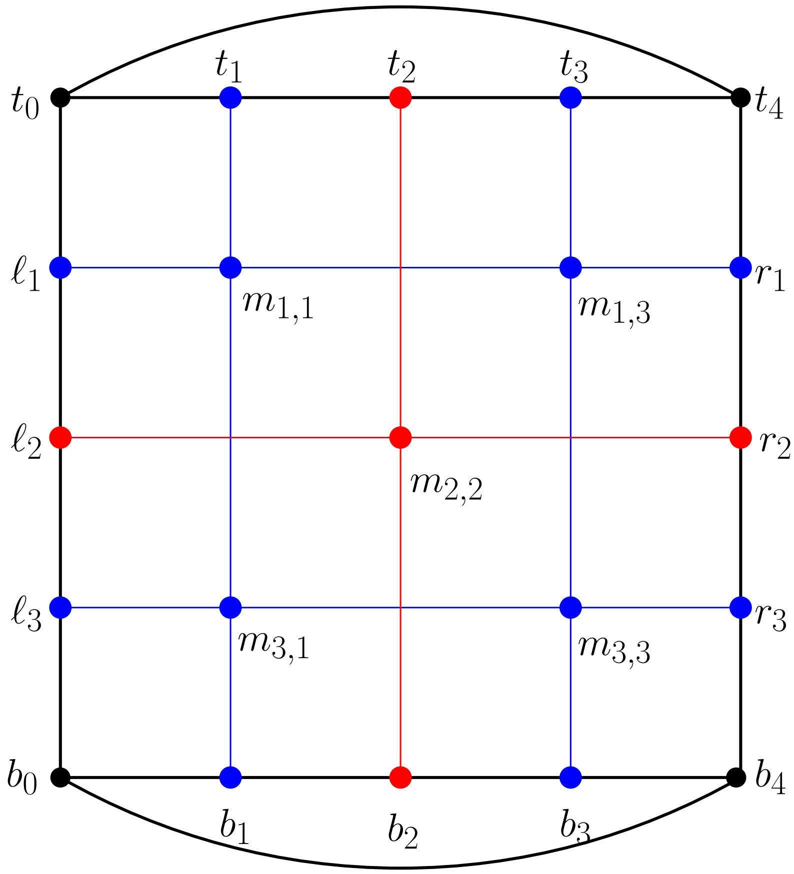

Figure 4.8: The new graph H_3^+

least 2 crossings by Lemma 4.3.5. So the original graph contained at least 3 crossings. A similar argument can be made if R_3 is crossed.

We will now assume that neither R_1 nor R_3 is crossed and show that no such optimal drawing can exist.

Case 1 *Suppose S is contractible* Let D be the disc bounded by S. If the cycle T is contractible and t_0t_4 is not contained in D, then each row and column is contained in D and there are at least 3 crossings. Then same can be said if B is contractible.

If both T and B are not contractible, they must be homotopic. Since S is contractible, L and R are contained in the same cylinder bounded by T and B. So the interior must be contained in this cylinder, and hence contained in the disc bounded by S.

As before, when each row and column is contained in the disc bounded by S, there are at least 3 crossings.

So if there is an optimal drawing of H_3^+ on the torus with T non-contractible, then B must be contractible. Moreover, if B is non-contractible then b_4b_0 must be drawn in D. Since B is uncrossed every column, and hence every row, must not be drawn in D. Then, t_4t_0 can be redrawn in D with no additional crossings.

So if there is an optimal drawing of H_3^+ on the torus with T non-contractible, there must be an optimal drawing with both t_4t_0 and b_4b_0 drawn in D.

Consider such a drawing. Let the cycle R_1^c be the cycle formed by R_1 together with the path $\ell_1t_0\dots t_4r_1$. Similarly, let R_3^c be the cycle formed by R_3 together with the path $\ell_3b_0\dots b_4r_3$. Note that if R_1^c is crossed, then R_1 is crossed since all other edges in R_1^c are thick. If R_3^c is crossed, then R_3 is crossed.

If either R_1^c or R_3^c is contractible, then c_2 must cross this cycle. However, we assumed R_1 and R_3 are uncrossed. If R_1^c and R_3^c are non-contractible, then they must be homotopic as otherwise each would be crossed.

If R_1^c and R_3^c are homotopic, they separate the torus into two cylinders. One cylinder must contain all thick edges in H_3^+ and hence the ends of R_2. The other cylinder contains the ends of all the columns, in particular c_2. Since c_2 and R_2 intersect, one of these must cross either R_1 or R_3. Hence, we can not have an optimal drawing with S contractible and R_1 and R_3 uncrossed.

Case 2 *Suppose S is non-contractible* We will examine this case in two possible sub cases depending on the ends of the row R_2.

Subcase (i) *Suppose the ends of R_2 are incident to the same side of S.* Either R_2 together with $\ell_1t_0\dots t_4r_1$ bounds a disc or R_2 together with $\ell_3b_0\dots b_4r_3$ bounds a disc. We will suppose it is the former, although the argument is symmetric for the other case. This implies that t_0t_4 cannot have ends in opposite sides of S, as otherwise it would cross R_2.

Now R_1 must have both ends on the same side of S, as otherwise it would cross R_2. If the ends of R_1 are on the same side as R_2, then R_1 and R_2 are homotopic and each column must cross the region bounded by them. We assume that R_1 has ends incident to side of S opposite of R_1. Now, there are two in $R_1 \cup S$ containing R_1, one must bound a disc. We will examine each of these two possibilities.

Suppose R_1 together with the path $t_0 \cdots t_4$ bounds a disc If t_4t_0 has ends in the same side of S as R_2, then c_2 crosses R_1. If the ends of t_4t_0 are in the opposite side, then c_1 must cross R_1. In either case, R_1 is crossed.

Alternatively, suppose R_1 together with the path $\ell_2\ell_3b_0 \dots b_4r_3r_2$ bounds a disc. So b_4b_0 cannot have ends in opposite sides of S.

If b_4b_0 has ends in the same side of S as R_1, then c_1 crosses R_1. So b_4b_0 has ends in the opposite side of S as R_1. If B is non-contractible, then either it crosses R_2 or $(S \setminus B) \cup \{b_0, b_4\}$ is a thick cycle which bounds a disc containing R_2. So either R_2 or c_2 must cross a thick edge. If B is contractible, then c_2 must cross R_1. We conclude that either R_1 or R_3 is crossed.

Subcase (ii) *Suppose the ends of R_2 are incident to opposite sides of S.* S and R_2 must be non-contractible and not homotopic.

If R_1 has ends on the same side of S, then it is contractible. Moreover, if R_1 is contractible, then t_4t_0 must have ends on the same side of S and must be contractible as well. If R_1 and t_4t_0 are incident to the same side of S, then c_1 must cross R_1. If R_1 and t_4t_0 are incident to opposite sides of S, then c_2 must cross R_1. Similarly, if R_3 is contractible, it must be crossed.

So we can assume that R_1 and R_3 are non-contractible and each have one end on each side of S. If we contract the path $t_4 \dots t_0Rb_0 \dots b_4$ to a point, we obtain a bouquet with four non-contractible loops. The rotation of this bouquet is $R_1, R_2, R_3, L, R_1, R_2, R_3, L$. By Theorem 3.2.1, two of these loops must be homotopic and due to the rotation the two homotopic loops must cross. Since L is thick and not crossed, either R_1 or R_3 is crossed.

We can conclude that R_1 or R_3 must be crossed in any optimal drawing of H_3^+ and hence, the toroidal crossing number is 3.

□

The above claim proves the existence of a graph with crossing sequence $(4, 3, 0)$. Moreover, the proof stipulates that every column is crossed on the torus. If we give column c_1 a weight of 2, we see that the crossing sequence becomes $(5, 3, 0)$. Thus, we have shown that for any $a > b > 0$, there exists a graph H with crossing sequence $(a, b, 0)$.

Chapter 5

The Graphs $G_{n,k}$

Theorem 1.0.1 proves the existence of a graph with an arbitrarily long, non-convex crossing sequence by specifying the sequence of a graph $G_{n,k}$. In particular, for any suitable n and k, the genus of $G_{n,k}$ is $2n+1$ and the ratio of 2 consecutive jumps in the crossing sequence is $1/(k-1)$. This chapter is devoted to describing the graph $G_{n,k}$ and finding its crossing sequence.

5.1 Description of $G_{n,k}$

To describe the graph $G_{n,k}$, let k and n be positive integers, with $k \geq 10$ even. Given n and k, let s be an arbitrary but fixed integer with $s \geq (3(2n+1)+2)^n + 1$.

Let T^0 be the graph obtained from a path of length s, replacing each edge with s parallel edges. The vertices in T^0 contain s edge disjoint paths of length s. We will label these paths $T_1^0, \dots, T_s^0$.

The vertices of T^0 are rigid and we will later use these parallel edges to prescribe the rotation of the vertices. Since the rotations are only prescribed up to inversion, a vertex must be incident to at least 4 edges in order for the rotation to be useful. However, in most other contexts it will suffice to consider T^0 simply a single path, and we will prove that this simplification does not affect the crossing number.

We similarly define s parallel paths $B_1^0, \dots, B_s^0$. Each of these parallel paths is also thick. We will often consider these paths as a single path B^0 of length s.

Now we define new paths $T^1, \ldots, T^{2n}$, each edge of which has weight $(nk)^2$, and such that T^j has length s if j is even and T^j has length k if j is odd. Let T be the concatenation of paths $T^0, T^1, \ldots, T^{2n}$. The path T is the *Top Path* and has length $nk + (n+1)s$.

Similarly, define the paths $B^1, \ldots, B^{2n}$, each edge of which has weight $(nk)^2$, and such that B^j has length s if j is even and B^j has length k if j is odd. Let the *Bottom Path* B be the concatenation of paths $B^0, B^1, \ldots, B^{2n}$, also of length .

Let the rigid vertex r be adjacent to the last vertex of both T and B. This new path of length 2 from T to B is the *Right Path*, denoted R and has weight $(nk)^2\ nk + (n+1)s$.

Now introduce a second rigid vertex ℓ, adjacent to the first vertex of T and B. Let $L_1, \ldots, L_s$ be parallel paths of length 2, each containing ℓ, and having the first vertex of both T and B as endpoints. Every edge of these parallel paths is also thick.

As with T^0 and B^0, we will use the parallel edges to fix the rotation of ℓ, but in most cases it will suffice to treat the parallel paths as a single path L, called the *Left Path*.

Note that $TLBR$ forms a cycle which we will call the *interior cycle*. Each edge in the interior cycle has weight at least $(nk)^2$.

Next, we introduce $n \cdot k$ paths of length 2. The *i, j-column* c^i_j is a path of length two between j^{th} vertex of T^{2i-1} and the j^{th} vertex of B^{2i-1} for each $1 \leq i \leq n$. Each column has weight nk.

We will often discuss the columns as if they are ordered from left to right. We say column $c^i_j > c^{i'}_{j'}$ iff $i > i'$ or $i = i'$ and $j > j'$.

In addition to the columns we define two rows, an *Even Row* R_e and an *Odd Row* R_o, each of weight 1. Both are paths of length $\frac{nk}{2}$ with endpoints ℓ and r. The path R_e is from ℓ to r with interior vertices consisting of the midpoints of the columns c^i_j, for all even j. The vertices are naturally ordered, as in T. Likewise, R_o is the path from ℓ to r with interior vertices consisting of the midpoints of the columns c^i_j, for all odd j. Again, the vertices of R_o are found in the natural order. Every vertex in R_e and R_o (including ℓ and r) are rigid vertices. We will often refer to the rows and columns together as the *interior* of the graph.

We note here that parity of k ensures that two consecutive columns are not incident to the same row. In fact, this is the only reason to dictate an even k. The proof for odd k is unchanged, however the description of the graph must be slightly altered to ensure that consecutive columns are not incident to the same row.

Finally, take $(3n+2)$ separate $s \times s$ grids with thick edges. Concatenate $2n+1$ of these grids by identifying the last column of one grid and the first column of the next. Attach

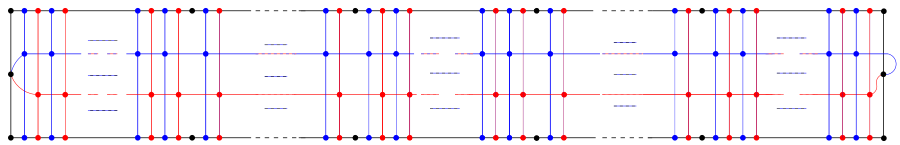

Figure 5.1: The interior of $G_{n,k}$ along with the interior cycle

the remaining $n+1$ grids evenly by identifying the top row of the grids with s consecutive vertices in the bottom row of concatenated grids. These result, as shown in Figure 5.2, is a gridlike structure resembling a dense "caterpillar" graph.

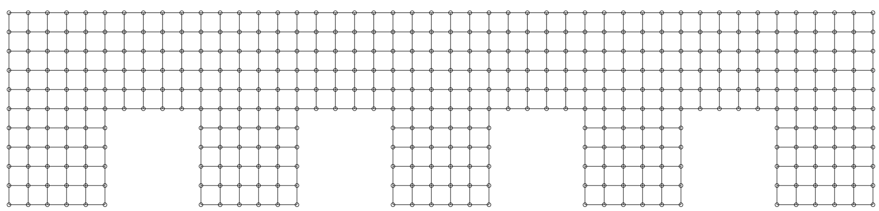

Figure 5.2: And example of the grid incident to the path T. (Here $n = 4$ and $s = 5$.)

The "base" of this gridlike structure contains $n+1$ sets of $s+1$ points. These points are identified with the points on the paths $T^0, T^2, \ldots, T^{2n}$.

Construct a symmetric gridlike structure for the Bottom Path. The "base" again contains $(n+1)$ sets of s points which are identified with the points on the paths $B^0, B^2, \ldots, B^{2n}$. Figure 5.3 shows a drawing of the interior cycle with the gridlike structures.

To complete the description of the graph, we must describe the rotation of the rigid vertices. For each path c^i_j, the rotation of the interior vertex is such that the column and incident row cross (and are not simply tangent to) one another.

At the vertex r, the two rows are incident to r on opposite sides of the path R.

At the vertex ℓ, the edges of $L_1, \ldots, L_s$ incident to T are found in the natural order followed by those incident to B in the reverse order. Both rows must now also be consecutive and are the only two edges between edges in L_1. In particular, the rows are incident to the same side of L. The rotations of ℓ and r are preserved in Figure 5.1.

Now we only need to consider the rotation of T^0 and B^0. Consider a planar embedding of the graph obtained by removing the rows from $G_{n,k}$. The rotations of T^0 and B^0 in G_{nk} are the same as rotations that result from such an embedding.

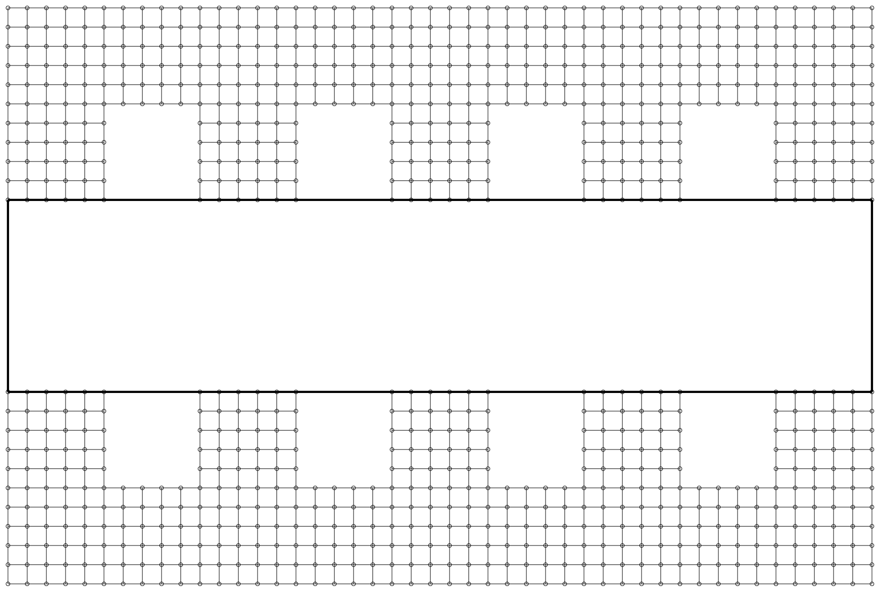

Figure 5.3: A drawing of $G_{n,k}$ without the interior. (Here $s = 5$.)

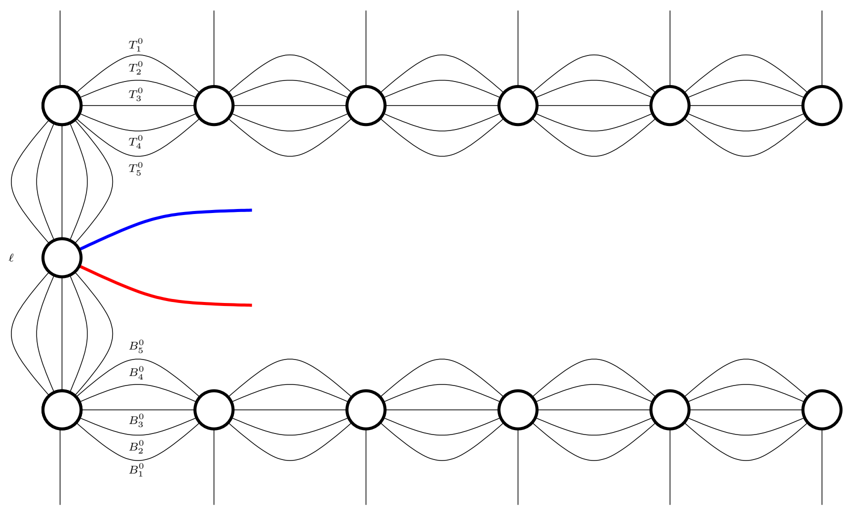

Figure 5.4: The paths T^0, L, and B^0 with the prescribed rotations.

5.1.1 Treating the Grids as Paths

The first application of Theorem 3.2.1 is in what we call the caps. In our description of the graph, we have dense gridlike structures, attached to the interior cycle at regular intervals. However, these grids are rather difficult to work with. Lemma 5.1.2 allows for much simpler treatment of these grids. To do this, we first must provide a bound for the crossing number on the torus.

Lemma 5.1.1. $\text{wcr}_1(G_{n,k}) < (nk)^2$.

Proof. If we combine Figures 5.1 and 5.3, we obtain a drawing on the plane with $2(nk)^2$ crossings. However, we see that one edge crosses both the last even column and the frame. By using one handle to reroute the this edge, we obtain a drawing of $G_{n,k}$ in the torus with fewer than $(nk)^2$ crossings. □

This bound on the toroidal crossing number will be used in the following lemma, since it ensures that the only crossings must include a row.

Lemma 5.1.2. *For any optimal drawing of $G_{n,k}$ on a surface of genus $1 \leq g \leq 2n+1$, there exists another drawing of $G_{n,k}$ on the surface of genus g with no additional crossings such that:*

(i) no edges in the grids are crossed; and

(ii) each edge in the grid incident to the interior cycle is incident to the same side of the interior cycle.

Proof. Consider an optimal drawing of $G_{n,k}$. We first wish to show that we need only consider drawings in which the parallel edges of T^0, B^0, and L are drawn as parallel arcs.

Consider two consecutive vertices u and v along the path $T^0 \cup L \cup B^0$. The two vertices are connected by s edges, which we will call $e_1, \ldots, e_s$, with $e_i \in T^0_i$. Moreover, the edges must be incident to both u and v in the natural order.

Contract e_s. The remaining $s-1$ edges form a RBG-bouquet and since $s - 1 \geq 3g$, by Theorem 3.2.1 at least two edges are homotopic or one is contractible. If two edges are homotopic, they were homotopic before contracting e_s. If an edge is contractible, then it was homotopic to e_s. So before contracting s, there were two edges e and e' that were homotopic. The edges e and e' must bound a disc in Σ.

Recall that the rotations of u and v can only be prescribed up to inversion. Assume that $e_1, \ldots, e_s$ are found in the natural order in the clockwise rotation of u. If $e_1, \ldots, e_s$ are also in the natural order in the clockwise rotation of v, then each edge would have exactly one end in the disc bounded by e and e'. So $e_1, \ldots, e_s$ must meet v in the reverse order.

Since e and e' are uncrossed, we can redraw the remaining edges $e_1, \ldots, e_s$ in a small neighborhood of e and e', while preserving the rotation and not increasing the crossing number.

So, given any optimal drawing, there is another optimal drawing in which the parallel paths of T^0, B^0 and L are drawn as parallel arcs. Moreover, since the grids are incident to T^0 and B^0 between the edges of T_1^0 and B_1^0, the grids must be incident to the same side of the path $T^0 \cup L \cup B^0$. We now need only consider optimal drawings in which the parallel edges are drawn in this way.

Define the path λ_j^i to be the path that begins at the j^{th} vertex at the base of the gridlike structure (i.e. the j^{th} vertex of T^0). The path travels up $s + (s - j)$ vertices, over $s \cdot (2i + 1) - 2j$ vertices, and then down $s + (s - j)$, ending at the $(s - j)^{th}$ vertex of T^{2i}.

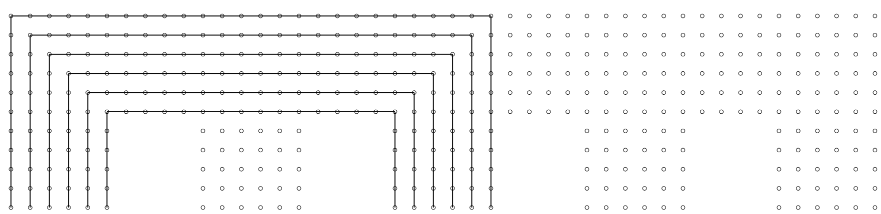

Figure 5.5: The paths $\lambda_1^2, \ldots, \lambda_s^2$.

For each i, the set of paths $\{\lambda_1^i, \ldots, \lambda_s^i\}$ forms a set of nested "arches" over a portion of the path T. Figure 5.5 shows an example of these paths. For a given j and any i_1, i_2, $\lambda_j^{i_1} \cap \lambda_j^{i_2}$ is non-empty.

Since the edges of the grids are thick, we know that none of them can be crossed, as otherwise the crossing number is not minimal.

We start by looking at the paths $\lambda_1^1, \ldots, \lambda_s^1$. Suppose we are walking along the path T. Each λ_j^1 will meet the path T twice. The first time we meet the path is along T^0. We know the grids are incident to the same side of T^0, which we assume is the left.

A path is said to be in D_1 if we find its ends on different sides of T, first to the left and then to the right. We say a path is in S_1 if we find it to the left on both occasions.

Let us now contract T to a point, so each path λ_j^1 becomes a loop incident to the point T. Note that the rotation of D_1 and S_1 is exactly as required for Lemma 3.2.1. Therefore, there can be at most $3g$ pairwise non-homotopic, non-contractible paths in $S_1 \cup D_1$. Including 2 possible contractible loops, there is at most $3g + 2$ non-homotopic loops in this bouquet.

Since $s \geq (3(2n+1)+2)^n + 1 \geq (3g+2)^n + 1$, the bouquet must contain at least $(3g+2)^n + 1$ loops. Moreover, since none of the paths is crossed, the rotation prohibits any path from D_1 from being homotopic to another path.

Therefore, S_1 must contain a set S_1' of at least $(3g+2)^{n-1}+1$ pairwise homotopic paths. Define the set $\Lambda_1 \subset \mathbb{N}$ by $j \in \Lambda_1$ if and only if $\lambda_j^1 \in S_1'$.

For each $1 < i \leq n$, we say a given path λ_j^i is in S_i if its endpoints are incident to the same side of T and recursively define the set Λ_i.

Let Λ_i to be those indices j for which $j \in \Lambda_{i-1}$ and $\lambda_j^i \in S_i$. Note that $\Lambda_1 \supseteq \Lambda_2 \supseteq \cdots \supseteq \Lambda_n$. We know that $|\Lambda_1| \geq (3g+2)^{n-1} + 1$. We wish to show that $|\Lambda_i| \geq (3g+2)^{n-i} + 1$, for each $i = 1, 2, \ldots n$.

By induction, $|\Lambda_{i-1}| \geq (3g+2)^{n-(i-1)} + 1$. Let $m = (3g+2)^{n-(i-1)}$ and let $j_1, j_2, \ldots, j_m$ be any m elements of Λ_{i-1}. By contracting T to a point and applying Theorem 3.2.1, we see that the m loops arising from $\lambda_{j_1}^i, \ldots, \lambda_{j_m}^i$ have at most $3g+2$ homotopy types.

By the pigeonhole principle, at least one homotopy class has size $(3g+2)^{n-i} + 1$. Since the paths are uncrossed, homotopic paths can not have ends on opposite sides of T. Therefore, the corresponding columns must be in Λ_i, as required.

In particular this tells us that $\Lambda_n \neq \emptyset$. Let $j \in \Lambda_n$. Then, $j \in \Lambda_i$, for each $1 \leq i \leq n$. By drawing the entire grid-like structure near $\lambda_j^1 \cup \cdots \cup \lambda_j^n$, we can modify our drawing of $G_{n,k}$ so for each $\ell \neq j$, the path λ_ℓ^i is homotopic to λ_j^i.

Since no λ_j^i is crossed, this gives a drawing of $G_{n,k}$ with no additional crossings. Moreover, each edge in the grid incident to T is incident to the same side of T, as required.

The argument is symmetric for the grids attached to the Bottom Path B. Since the edges of the grid are incident to the same side T^0LB^0, we can conclude that this new optimal drawing has all edges of the grid incident to the interior cycle are incident to the same side of the interior cycle. □

Lemma 5.1.2 shows that we only need to consider drawings where the set of paths $\{\lambda_1^i, \ldots, \lambda_s^i\}$ are homotopic, for any $1 \leq i \leq n$. Equivalently, our gridlike structure has n "arches", and the homotopy of these arches defines the homotopy of the rest of the structure.

For the grid attached to the Top Path, let α_i be the innermost path around the i^{th} arch, so in particular it joins the end of T^{2i-2} with the beginning of T^{2i}. For the bottom grid, let β_i be the innermost path around the i^{th} arch, joining the end of B^{2i-2} with the beginning of B^{2i}. Figure 5.6 shows the paths $\alpha_1, \ldots, \alpha_n$.

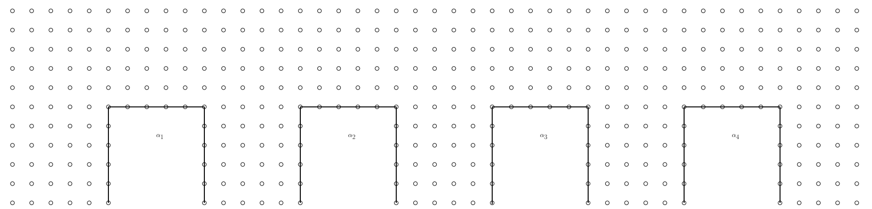

Figure 5.6: A drawing of the caps $\alpha_1, \ldots, \alpha_n$.

We will call the paths α_i and β_i the *caps*. The cap α_i along with the path T^{2i-1} forms a cycle which we will call the i^{th} *upper cap cycle*, denoted $\triangle_i$. Similarly the cycle formed by β_i and B^i is a *lower cap cycle*, denoted $\bigtriangledown_i$. Finally, the paths L and R together with the paths $\alpha_i, \beta_i, T^{2i}$ and B^{2i}, for $0 \leq i \leq n$, form a cycle, which we will call the *exterior cycle*. Together, the interior and exterior cycles form the *frame*.

Figure 5.7 shows a complete drawing of the graph $G_{n,k}$. We assume that the remainder to the grid like structures are drawn close to the cap paths α_i and β_i.

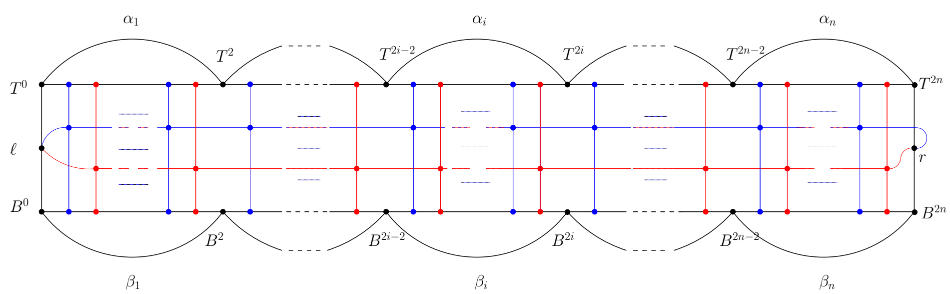

Figure 5.7: The graph $G_{n,k}$.

5.2 Using Homotopy to Analyze $G_{n,k}$

In general, we will be using the standard definitions of homotopy when we talk about the cycles and paths in our graph. However, in our discussion of the homotopy of the columns of $G_{n,k}$, we will use the same definition that we saw in the discussion of the paper of DeVos et al. We restate it here.

Two columns, $c^i_j > c^{i'}_{j'}$, are *homotopic* if the cycle contained in $T \cup B \cup c^i_j \cup c^{i'}_{j'}$ bounds a closed disc. (In this case, we will simply say c^i_j and $c^{i'}_{j'}$ bound a disc.)

A column, c_j^i is *contractible* if the cycle contained in $T \cup B \cup L \cup c_j^i$ or the cycle contained in $T \cup B \cup R \cup c_j^i$ bounds a closed disc. We will call c_j^i *left contractible* former case and *right contractible* in the latter.

As in the paper of DeVos et al., we will need that if any set of j consecutive columns are homotopic, then they must contain at least j crossings. The statement of the following lemma is analogous to Lemma 4.3.2 and only requires small adjustments to the proof.

Lemma 5.2.1. *Suppose $G_{n,k}$ is drawn optimally on a surface Σ. If there are $j \geq 2$ consecutive columns, $c_1, \dots, c_j$ that are pairwise homotopic, then these columns are crossed by the two rows at least j times in total.*

Proof. Suppose $j = 2$. Since c_1 and c_2 are homotopic, they bound a disc. By Lemma 5.1.2, the exterior cycle must be disjoint from the disc bounded by c_1 and c_2, as otherwise the crossing number would be at least $(nk)^3$. Therefore, ℓ and r are either both inside this disc or both outside.

Each row has a vertex in the boundary of the this disc bounded by c_1 and c_2. The rotation of this vertex is prescribed in such a way that the row must cross the boundary of the disc at the vertex. Therefore, it must cross the boundary of the disc at another point. Thus, both rows cross the boundary of the disc, as required.

Suppose $j = 3$ and that c_1, c_2, and c_3 are consecutive homotopic columns. Then c_2 is contained in the disc bounded by c_1 and c_3. Then the row that contains the vertex in c_2 must cross $c_1 \cup c_3$ twice. The other row must cross $c_1 \cup c_2$ and we are done.

Suppose now that $j \geq 4$ and let $c_1, \dots, c_j$ be consecutive homotopic columns. By induction, $c_1, \dots, c_{j-2}$ contains $j - 2$ crossings. Additionally, each row crosses the disc bounded by c_{j-1} and c_j, completing the proof. □

A further application of Theorem 3.2.1 is to count the homotopy types of the columns. Once the number of homotopy types is bounded, we will be able to use Lemma 5.2.1 to obtain a lower bound on the number of crossings.

We will see in Lemma 5.4.4 that there are two ways to save crossings in our graph. The first is to make both caps in a pair non-contractible. If this is the case for a given drawing, we say that $\{\triangle_i, \triangledown_i\}$ is an *active pair*. If a pair of caps $\{\triangle_i, \triangledown_i\}$ is not active, we call it an *inactive pair*. Generally, if $\triangle_i$ is in an active pair, all the columns under $\triangle_i$ are uncrossed.

The second way to save crossings is by maximizing the homotopy classes of the columns. The following lemma provides an upper bound for the number of homotopy types for a given surface.

Theorem 5.2.2. *Suppose we have a drawing $G_{n,k}$ on a surface of genus g in which ℓ columns are all under inactive pairs, pairwise non-homotopic, and do not cross the frame. Then, $\ell \leq 3g + 1$.*

Proof. As we walk along the path TLB, we will meet each of the ℓ columns exactly twice. Assuming, without loss of generality, that we find the path α_n to our right, there are three possible scenarios in which we can meet any given column, which we will think of as a partition of the columns.

We say that a column is in R if we find find it to our right the first time we meet it and to our left the second time. We say that a column is in B if we find find it to our left the first time we meet it and also to our left the second time. We say that a column is in G if we find find it to our left the first time we meet it and to our right the second time.

Let us define $r_1, \ldots, r_i$ to be the columns in R where $r_m = c_{m_2}^{m_1}$ precedes $r_n = c_{n_2}^{n_1}$ iff $m_1 < n_1$ or $m_1 = n_1$ and $m_2 < n_2$. We similarly define $g_1, \ldots, g_j$ and $b_1, \ldots, b_k$ to be the sets of columns in G and B, respectively.

Now let us contract the path TLB to a point, s. We are left with a bouquet with ℓ loops in a particular rotation. This rotation is the same as required by Theorem 3.2.1. Then, taking into account the possibility of a contractible column, $\ell \leq 3g + 1$. □

This bound is important because if two consecutive columns are homotopic, then by Lemma 5.2.1 they contribute at least two crossings.

The last step before we analyze the crossing number of our graph on surfaces of higher genus is to limit the ways in which the interior cycle and caps can be drawn on the surface. We will use the following propositions several times in our analysis.

Proposition 5.2.3. *Let $g \geq 1$ and $\mathcal{D}$ be an optimal drawing of $G_{n,k}$ on the surface of genus g. Then, the interior cycle is non-separating in $\mathcal{D}$.*

Proof. Suppose the interior cycle is separating in some drawing of $G_{n,k}$. Then the interior cycle separates the surface into 2 regions. Due to the rotation of ℓ and r, one row must be incident to ℓ in one region and incident to r in the other region. This row must cross the boundary interior cycle, so the drawing has at least $(nk)^2$ crossings. By Lemma 5.1.1, this drawing is not optimal. □

Proposition 5.2.4. *Let $g \geq 1$ and $\mathcal{D}$ be an optimal drawing of $G_{n,k}$ on the surface of genus g. Then, the interior cycle is not homotopic to any non-contractible cap in $\mathcal{D}$.*

Proof. From Proposition 5.2.3, we know that the interior cycle is non-contractible in $\mathcal{D}$. Suppose α_i is a non-contractible cap. If it is homotopic to the interior cycle, then together the two curves bound a disc D. The vertices ℓ and r are on the boundary of this disc.

Due to the rotation of ℓ and r, one row must be incident to exactly one of ℓ or r in D. This row must cross the boundary of D, so the drawing has at least $(nk)^2$ crossings. By Lemma 5.1.1, $\mathcal{D}$ is not optimal, a contradiction. □

Finally, our analysis will eventually reduce to the case of the double torus. While Theorem 5.2.2 provides a bound for the number of homotopy types on the double torus, we are able to find a tighter bound on the number of crossings.

Lemma 5.2.5. *Suppose we have an optimal drawing $G_{n,k}$ on the double torus in which all columns are all under inactive pairs. Then there are at least $nk - 2$ crossings of columns.*

Proof. In the double torus, by Proposition 5.2.3 we have that the interior cycle must be non-separating. Moreover, since each edge in the interior cycle has eight at least $(nk)^2$, we know from Lemma 5.1.1 that the cycle is uncrossed.

Case 1 *All caps are contractible.* In this case, the interior cycle must be homotopic to the exterior cycle. If we cut and cap along the exterior cycle, then there are two copies of the exterior cycle in the torus, each bounding a disc. Since the interior cycle is homotopic to the exterior cycle, it must also bound a disc in the torus.

Since the interior cycle is uncrossed and, by Lemma 5.1.2(i), the exterior cycle is as well, the interior of $G_{n,k}$ is drawn outside this disc in an optimal drawing. In particular, both ends of each column are incident to the same side of the interior cycle. We will contract the interior cycle and let the columns form a bouquet on the torus.

We know that the rotation of the columns must now be $c_1, c_2, \ldots, c_{nk}$ in the first half and $c_{nk}, \ldots, c_2, c_1$ in the second half. By Corollary 3.2.2 there can only be one non-contractible loop on the torus in such a bouquet. If there is a non-contractible loop, let it be c_i.

The remaining columns must be either left or right contractible. Suppose c_j is left contractible. Then, for all $k \leq j$, the disc bounded by c_j contains both ends of the column c_k. Since two columns cannot cross each other, c_k must be left contractible as well. Similarly, if $c_{j'}$ is right contractible, then $c_{k'}$ is right contractible, for all $k' \geq j'$.

Let j be the largest integer such that c_j is left contractible. If $j \leq 1$, then c_3 is right contractible (c_2 may be right contractible or non-contractible). So $c_3, \ldots c_{nk}$ are

$nk-2$ consecutive homotopic columns and contain $nk-2$ crossings by Lemma 5.2.1. If $j \geq nk-2$, then $c_1, \ldots, c_{nk-2}$ are homotopic and likewise contain $nk-2$ crossings.

If $2 \leq j \leq nk-3$, then $c_1, \ldots, c_j$ are left contractible and $c_{j+2}, \ldots, c_{nk}$ are right contractible. By Lemma 5.2.1, the left contractible columns contribute j crossings and the right contractible columns contribute $nk-j-1$ for a total of $nk-1$. So we see that the columns are crossed at least $nk-2$ times.

Case 2 *Some cap is non-contractible.* In this case, we first note that the non-contractible cannot be separating. If it were separating, the remainder of the interior cycle must be drawn in one region, since the cap is uncrossed in any optimal drawing. Then the remaining caps and the interior of the graph are drawn in the same region as the interior cycle. If we replace the empty face bounded by the original cap with a disc, we obtain a drawing with the same number of crossings but in a surface of smaller genus. This contradicts the optimality of the drawing.

So the non-contractible cap must be non-separating. Moreover, if it were homotopic to the interior cycle, the cap together with the interior cycle would bound a disc. Due to the rotation of r and ℓ, one row would have exactly one end inside this disc and would cross the frame.

Assuming that the non-contractible cap is $\triangle_i$, let us cut along it. Since the interior cycle was not homotopic to $\triangle_i$, it must be non-contractible in the new surface. Cut along the interior cycle as well. We are left with a sphere with four holes, but it is important to note that we know the boundaries of these discs.

One hole, h_1 is bounded by $\triangle_i$, with the points of T^i appearing in clockwise order.

A second disc h_2 is bounded by the interior cycle. Since Σ was an orientable surface, the two faces containing T^i after cutting and capping will have the vertices of T^i appearing in opposite directions. So the vertices of T^i appear in counter-clockwise order around h_2. The vertices of B^i must then appear in clockwise order.

The final holes h_3 and h_4 are bounded by $\triangle_i$ and the interior cycle, respectively. However, since these two cycles are not disjoint, the boundary of these holes intersect at the path T^i.

Since all the caps are incident to the same side of the interior cycle, all other caps are incident to the interior cycle on the boundary of h_4. So the exterior cycle bounds a region containing the holes h_3 and h_4. This means that all columns either have both ends incident to h_2 or one end in h_1 and one in h_2.

Suppose there are at least two columns having one end on the boundary of h_1 and one end on the boundary of h_2. Then they necessarily separate the sphere. Since the other copy of the interior cycle (namely the boundary of h_4) is contained in a region bounded by the exterior cycle, these two columns together with the exterior cycle separate the double torus.

Each row intersects one of the two columns at exactly one vertex. Due to the prescribed rotation, it crosses the boundary of the separation at this vertex. Since ℓ and r are in the same region of this separation, each row must therefore cross the boundary at another point. So, these two columns are crossed by each row.

Moreover, since the directions of T^i in h_1 and B^i in h_2 are the same, the columns must be consecutive. In particular, there can not be more than two of them, so every other column has both ends in h_2.

Each column with both ends in h_2 must be either left or right contractible in the double torus. As we stated in Case 1, if a column is left contractible, then all preceding columns are left contractible. If a column is right contractible, then all successive columns are right contractible.

If every column has both ends in the boundary of h_2, then every column is either left or right contractible. By Lemma 5.2.1, there are at least $nk - 1$ crossings, achieved when one of the contractible homotopy classes contains a single column.

If exactly one column has one end in the boundary of h_1 and the other in the boundary of h_2, then there are at least $nk - 2$ crossings. Again, this is achieved when one of the contractible homotopy classes contains 1 column. The other class now contains $nk - 2$ columns and, by Lemma 5.2.1, they are crossed $nk - 2$ times.

Finally, if two columns have one in h_1 and the other end in h_2, then they are crossed by each row. The remaining $nk - 2$ columns must contain $nk - 3$ crossings, again achieved when one column is in its own homotopy class and the remaining $nk - 3$ columns are in the other. In total, this scenario contains $nk - 1$ crossings.

Thus, it is optimal to have only one such column with one end in h_1 and the other in h_2 and there are at least $nk - 2$ crossings. □

5.3 Crossing Number on the Sphere and Torus

Before we get into a more general case, we begin by determining $\mathrm{cr}_0(G_{n,k})$ and $\mathrm{cr}_1(G_{n,k})$.

Lemma 5.3.1. $\mathrm{wcr}_0(G_{n,k}) = 2n^2k^2$

Proof. Figure 5.7 gives a drawing in the plane with exactly $2n^2k^2$ crossings.

Consider an optimal drawing $\mathcal{D}$ of $G_{n,k}$ in the sphere. The interior cycle must separate the sphere into two discs, d_1 and d_2. Since ℓ and r have fixed rotation, one of the rows must cross the interior cycle. This contributes n^2k^2 to the crossing number.

Recall that each edge in the interior cycle has weight at least $(nk)^2$ and each column has weight nk. No column can cross the interior cycle, as otherwise the crossing number would be at least $(nk)^3$. Suppose a column c_j^i is in d_1. Since $\triangle_i$ and $\triangledown_i$ are thick and uncrossed, they must both be in d_2. Then, $c_1^i, \ldots, c_k^i$ must all be in d_1.

So for each i, the columns $c_1^i, \ldots, c_k^i$ are homotopic and, by Lemma 5.2.1, are crossed k times by the rows. So the nk columns combined are crossed nk times. Since each column has weight nk, these contribute $(nk)^2$ to the crossing number. Therefore, the crossing number is at least $2(nk)^2$ and equality holds. $\square$

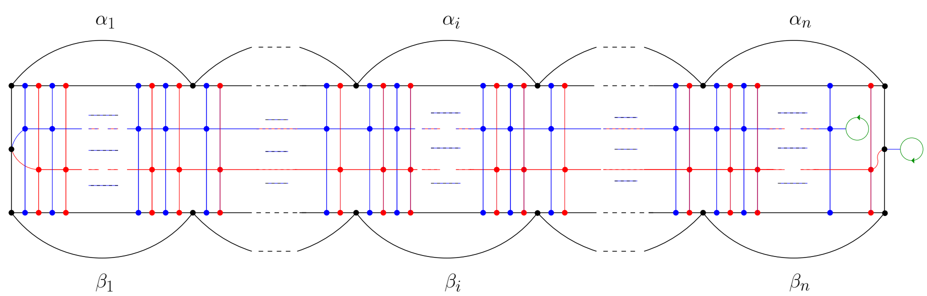

Figure 5.8: Drawing on torus with $n^2k^2 - nk$ crossings

Lemma 5.3.2. $\mathrm{wcr}_1(G_{n,k}) = n^2k^2 - nk$.

Proof. Figure 5.8 provides us with an upper bound of $\mathrm{cr}_1(G_{n,k}) \leq n^2k^2 - nk$.

To bound from below, we begin by noting that the frame cannot be crossed, as otherwise there would be at least $(nk)^2$ crossings. The prescribed rotations at ℓ and r imply that the interior cycle must be non-separating.

If a cap $\triangle_i$ is non-contractible, then we can contract the path T^{2i-1}, so that $\triangle_i$ and the interior cycle form a bouquet of non-contractible loops on the torus. Since the ends of the caps are incident to the same side of the interior cycle, they must be consecutive after contracting T^{2i-1}. By Corollary 3.2.2, we see that these two loops must be homotopic.

If a cap is homotopic to the interior cycle then together, since the two cycles share a path, they bound a disc in the torus. Due to the rotation of ℓ and r, one row would have exactly one end that enters the interior of this disc. This row would cross the boundary of the disc and contribute at least $(nk)^2$ crossings. Since the drawing was optimal, no such non-contractible cap can exist.

Now, we can cut and cap along the exterior cycle; this will not separate the interior cycle. By Lemma 5.1.2(i), we see that the columns are now drawn on the sphere and must be either left or right contractible. In particular, the first j columns must be left contractible and last $nk - j$ columns must be right contractible.

Lemma 5.2.1 implies that there are at least $nk-1$ crossings of columns, where equality can only hold if $j = 1$ or $nk-1$. Since each column has weight nk, we conclude $\text{cr}_1(G_{n,k}) = n^2k^2 - nk$. □

5.4 $G_{n,k}$ on Surfaces of Higher Genus

We now turn our attention to the drawing of $G_{n,k}$ on a surface of arbitrary genus. To do this, we separate into cases when the genus is even or odd. We will then find upper and lower bounds for the crossing number in each case.

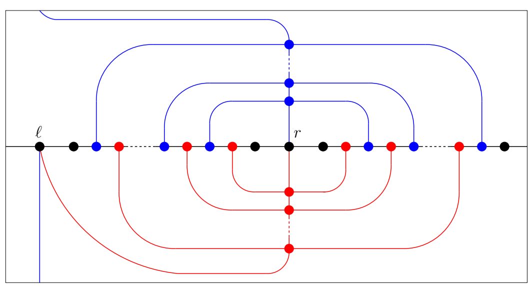

Figure 5.9: An embedding of the interior of $G_{n,k}$ on the torus

Lemma 5.4.1. *$G_{n,k}$ embeds in the surface of genus $2n+1$.*

Proof. As demonstrated by Figure 5.9, the interior of the graph embeds in the torus.

Consider the subgraph of $G_{n,k}$ consisting of the interior and j caps and assume the subgraph is embedded on the surface of genus $j+1$. We will show that we can embed a subgraph consisting of the interior and $j+1$ caps on the surface of genus $j+2$.

Suppose the cap path α_i is not in our subgraph embedded in the surface of genus $j+1$. The argument for β_i is symmetric. Let f_i be a face with boundary containing the path

T^{2i-1} and f_i' be a face containing the path T^{2i+1}. If we attach a handle with an end in f_1 and f_1', then the cap α_i can be drawn without introducing any more crossings. This yields an embedding of the interior and $j+1$ caps on the surface of genus $j+2$.

By induction, we conclude that there is an embedding of the interior will all $2n$ caps on the surface of genus $2n+1$. We now only need to specify how to draw the remainder of the gridlike structure.

First consider the top gridlike structure. For $\lambda_1^1, \ldots, \lambda_s^1$, draw each path alongside α_1. Now suppose we have drawn $\lambda_1^{i-1}, \ldots, \lambda_s^{i-1}$. We draw the remaining portion of $\lambda_1^i, \ldots, \lambda_s^i$ alongside α_i. This specifies the full grid. We can embed the opposite grid like structure in a similar way to obtain an embedding of $G_{n,k}$ in the surface of genus $2n+1$. □

Lemma 5.4.2. *Let $1 \leq g \leq n$, then:*

(a) $\operatorname{wcr}_{2g-1}(G_{n,k}) \leq [nk-(g-1)k-1]nk$; *and*

(b) $\operatorname{wcr}_{2g}(G_{n,k}) \leq [nk-(g-1)k-2]nk$

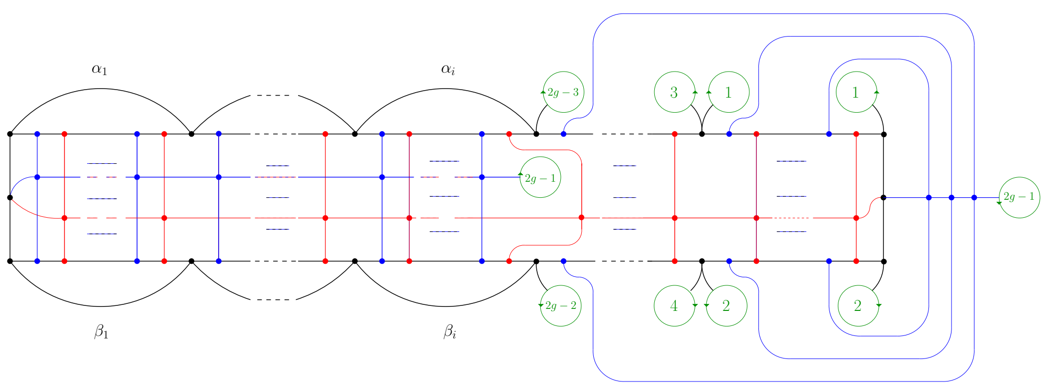

Figure 5.10: A drawing of $G_{n,k}$ on a surface of genus $2g-1$

Proof. For the case of $2g-1$, this upper bound is realized in Figure 5.10. The first $2g-2$ handles are used to make $g-1$ active pairs. The final handle is used to make the interior cycle non-contractible.

To achieve the bound for the surface of genus $2g-1$, begin with the drawing on the surface of genus $2g-1$ in Figure 5.10. Note that every crossing is a row (of weight 1) on a column (of weight nk). In particular, the row R_e crosses at least one column. Choose one such crossing. We can introduce a handle to locally remove this crossing. The result is a drawing of $G_{n,k}$ on the surface of genus $2g$ with $[nk-(g-1)k-2]nk$ crossings. $\square$

Before we prove that this bound is tight, we would like to limit the types of drawings of $G_{n,k}$ that we must consider. The following lemma shows that we need only consider drawings of $G_{n,k}$ in which no subset of the non-contractible cap cycles separate the surface.

Lemma 5.4.3. *Let $\mathcal{D}$ be an optimal drawing of $G_{n,k}$ on a surface of genus $1 \le g \le 2n$. Then there exists another optimal drawing $\mathcal{D}'$ of $G_{n,k}$ on a surface Σ of genus g such that no subset of non-contractible cap cycles separates Σ.*

Proof. We first note that every edge in a cap cycle has weight at least $(nk)^2$ and is uncrossed in the optimal drawing of $G_{n,k}$. Let $\mathcal{C}$ be a smallest subset of non-contractible cap cycles that separates the surface. Consider the component of $G_{n,k} \setminus \mathcal{C}$ containing ℓ. Since $\mathcal{C}$ is uncrossed, this component must be contained in the same region of $\Sigma \setminus \mathcal{C}$. Moreover, this component contains all of $G_{n,k}$ except for $\mathcal{C}$ and the paths T^{2i} and B^{2j}, for $\triangle_i, \triangle_{i+1} \in \mathcal{C}$ and $\triangledown_j, \triangledown_{j+1} \in \mathcal{C}$ respectively.

Consecutively cut and cap along each curve in $\mathcal{C}$. By the minimality of $\mathcal{C}$, the first $|\mathcal{C}|-1$ curves are non-separating, so cutting them yields a surface of genus $g-|\mathcal{C}|+1$. Cutting the final curve yields two surfaces Σ_1 and Σ_2 of genus g_1 and g_2 respectively, with $g_1+g_2 = g-|\mathcal{C}|+1$.

For each curve $c \in \mathcal{C}$, there is a copy of c drawn in Σ_1 and Σ_2, with both copies of c contractible.

We will assume that Σ_1 contains ℓ, and thus must contain all of $G_{n,k}$ with the possible exception of the paths T^{2i} and B^{2j}, for $\triangle_i, \triangle_{i+1} \in \mathcal{C}$ and $\triangledown_j, \triangledown_{j+1} \in \mathcal{C}$ respectively. However, there can be at most $|\mathcal{C}|$ such paths. Let T_{2i} be one such path that is not drawn in Σ_1.

Add a handle to Σ_1 with one end in a face with boundary containing the end of the path T^{2i-1} and other end in a face with boundary containing the beginning of the path T^{2i+1}. Then, the path T^{2i} can be drawn in Σ_1 (with an additional handle) without introducing any new crossings.

This can be repeated for every such path T^{2i} or B^{2j} until we are left with a complete drawing of $G_{n,k}$ in Σ_1 with at most $|\mathcal{C}|-1$ additional handles. Since $g_1 \le g-|\mathcal{C}|+1$, we see that the new drawing is on a surface of genus no greater than g. Since the original drawing was optimal, we must have equality.

So we now having an optimal drawing of $G_{n,k}$ in the surface of genus g with each cap in c now contractible. We can repeat this process until no subset of non-contractible cap cycles separates the surface. □

The following lemmas will complete our analysis of the crossing numbers of $G_{n,k}$.

Lemma 5.4.4. *If $1 \leq g \leq n$, then:*

$$\text{(a)} \quad \text{wcr}_{2g-1}(G_{n,k}) = [nk - (g-1)k - 1]nk; \textit{ and}$$

$$\text{(b)} \quad \text{wcr}_{2g}(G_{n,k}) = [nk - (g-1)k - 2]nk$$

Proof. Consider an optimal drawing of $G_{n,k}$ on a surface of genus $g' \geq 1$ and suppose that there are a active pairs. (Here g' will be $2g$ or $2g-1$, but at the moment we do not wish to consider the parity of the genus).

By Lemma 5.4.3, these $2a$ caps do not separate the surface. Successively cut and cap along each of the caps to obtain a surface of genus $g' - 2a$. If this new surface is the sphere, then the interior cycle separates the surface and all columns must be drawn in one disc. By Lemma 5.2.1, the columns are crossed at least nk times. If this were the case, the drawing is not optimal. We conclude that the new surface is not the sphere, so $g' - 2a > 0$. In particular, note that $a \leq g - 1$.

Let $X(g', n)$ be the minimum number of crossings when $G_{n,k}$ is drawn on the surface of genus g' with no active pairs. Since there are no active pairs, if two columns c_i and c_k are homotopic, with $i < j < k$, c_j must have at least one end in the disc bounded by c_i and c_k. Since columns do not cross, c_j must also be homotopic to c_i and c_k.

So if there are $j \geq 2$ homotopic columns, they must be consecutive and, by Lemma 5.2.1, must be crossed by the rows j times. By Theorem 5.2.2, there are at most $3g'+1$ homotopy classes on the surface of genus g'. So at least $nk - 3g'$ columns are not in singleton homotopy classes. We see that $X(g', n) \geq nk[nk - 3g']$.

Again, we consider an optimal drawing of $G_{n,k}$ with a active pairs. We cut along each cap in an active pair. For each active pair $\{\triangle_i, \triangledown_i\}$, we will assume all columns $c_1^i, \ldots, c_k^i$ are uncrossed and remove them from the drawing. Further, we contract $T^{2i-1}, T^{2i}, B^{2i-1}$, and B^{2i}. Finally, remove the appropriate edges of the grid like structures so that we are left with a drawing of $G_{n-a,k}$ on the surface of genus $g' - 2a$.

The drawing of $G_{n-a,k}$ has no active pairs and, consequently, at least $X(g' - 2a, n - a)$ crossings. Therefore, so does the original drawing of $G_{n,k}$ with a active pairs.

It is clear that $\text{wcr}_{g'}(G_{n,k}) \geq \min_{0 \leq a \leq g-1} X(g'-2a, n-a)$. Since $k \geq 6$, we see that the minimum of $X(g'-2a, n-a)$ occurs when a is large. Further, we know $1 \leq a \leq g-1$. When $a = g-1$, the remaining surface is either the torus or double torus. In either case, we have a tighter bound for the crossing number. We look at these two cases separately and then compare it to the general bound obtained from $X(g'-2a, n-a)$ when $a = g-2$.

Case 1 *Suppose the original genus is* $2g$. If $a = g-1$, then after cutting and capping, we are left with a drawing of $G_{n-g+1,k}$ on the double torus. By Lemma 5.2.5, we see that the crossing number is at least $[nk-(g-1)k-2]nk$

We see that $X(2g-2a, n-a) \geq nk[nk-(g-2)k-12]$, so $\text{wcr}_{2g}(G_{n,k}) \geq \min\{[nk-(g-1)k-2]nk, [nk-(g-2)k-12]nk\}$. Since $k \geq 10$, $\text{wcr}_{2g} \leq [nk-(g-1)k-2]nk$ and by Lemma 5.4.2(a) equality holds.

Case 2 *Suppose the original genus is* $2g-1$. If $a = g-1$, after cutting and capping, we are left with a drawing of $G_{n-g+1,k}$ on the torus. By Lemma 5.3.2, we see that the crossing number is at least $[nk-(g-1)k-1]nk$.

We see that $X(2g-1-2a, n-a) \geq nk[nk-(g-2)k-9]$, so $\text{wcr}_{2g}(G_{n,k}) \geq \min\{[nk-(g-1)k-1]nk, [nk-(g-2)k-9]nk\}$. Since $k \geq 10$, $\text{wcr}_{2g}(G_{n,k}) \leq [nk-(g-1)k-1]nk$ and by Lemma 5.4.2(b) equality holds.

□

Lemma 5.3.1 ($g = 0$), Lemma 5.4.2 ($1 \leq g \leq 2n$), and Lemma 5.4.1 ($g = 2n+1$) give the entire crossing sequence for $G_{n,k}$, summarized in the following statement.

Theorem 5.4.5. *Let n be a positive integer and $k \geq 10$ be an even positive integer. Then, the crossing sequence for $G_{n,k}$ is given by:*

$$\begin{aligned}
\text{wcr}_0(G_{n.k}) &= 2n^2k^2 \\
\text{wcr}_{2g-1}(G_{n,k}) &= [nk-(g-1)k-1]nk \quad \textit{for } 1 \leq g \leq n \\
\text{wcr}_{2g}(G_{n,k}) &= [nk-(g-1)k-2]nk \quad \textit{for } 1 \leq g \leq n \\
\text{wcr}_{2n+1}(G_{n.k}) &= 0
\end{aligned}$$

From the discussion of special graphs in Section 2.2 and Section 2.3, there must exist a simple graph with the same (unweighted) crossing sequence, proving the main result, Theorem 1.0.1.

Theorem 5.4.5 gives the entire crossing sequence of $G_{n,k}$. Some interesting properties of the sequence include:

(a) a non-convex jump can occur arbitrarily late in the sequence;

(b) this late non-convex jump can be arbitrarily large; and

(c) there can be arbitrarily many non-convex jumps.

Chapter 6

Conclusion

The main result of this paper proves the existence of a graph with an arbitrarily long non-convex crossing sequence. In particular, the convex jumps occur throughout the sequence, so for any $N \in \mathbb{N}$, there exists an $n \geq N$ such that $\mathrm{cr}_{n-1} - \mathrm{cr}_n < \mathrm{cr}_n - \mathrm{cr}_{n+1}$. Moreover, for any such n, the size of this non-convex jump remains arbitrarily large.

In [1], Archdeacon et al. make following conjecture:

Conjecture 6.0.6. *Any strictly decreasing sequence of non-negative integers is the crossing sequence of some graph.*

Though our result provides further evidence towards this conjecture, progress must still be made in order to achieve a full proof. In particular, our graph was the first known example to contain multiple non-convex jumps. However, the non-convex jumps remained constant in size throughout the length of the sequence. One goal for future research would be to prove the existence of a crossing sequence with multiple non-convex jumps in which the sizes of the jumps increase.

Additionally, the non-convex jumps in the crossing sequence of $G_{n,k}$ are spaced evenly throughout the sequence. Another possible direction for research would be to create a family of graphs where the non-convex jumps in the crossing sequence can be controlled through a parameter of the graph. One strategy to achieve this result may be to relax any restriction regarding the relative size of the non-convex jumps.

References

[1] D. Archdeacon, C. P. Bonnington, and J. Širáň. Trading crossings for handles and crosscaps. *J. Graph Theory*, 38:230–243, 2001.

[2] J. Battle, F. Harary, Y. Kodama, and J. W. T. Youngs. Additivity of the genus of a graph. *Bull. Amer. Math. Soc.*, 68:565–568, 1962.

[3] M. DeVos, B. Mohar, and R. Šámal. Unexpected behaviour of crossing sequences. *J. Combin. Theory, Ser. B*, 101(6):448–463, 2011.

[4] B. Mohar and C. Thomassen. *Graphs on Surfaces.* The Johns Hopkins University Press, Baltimore, Maryland, 2001.

[5] J. Stillwell. *Classical Topology and Combinatorial Group Theory.* Springer, New York, 2003.

[6] J. Širáň. The crossing function of a graph. *Abh. Math. Sem. Univ.*, 53:131–133, 1983.