Freie Universität Berlin
Fachbereich Mathematik und Informatik

Gauss Codes and Thrackles

On Characterizations of Closed Curves in the Plane with an Application to the Thrackle Conjecture

Felix Breuer

Diplomarbeit

Betreuer: Prof. Dr. Martin Aigner
Dr. Mark de Longueville

29. September 2006

Table of contents

Introduction

Our topic are closed curves in the plane that have finitely many double points and no other multi points. With every such curve $c\colon S^1 \to \mathbb{R}^2$ we can associate a multigraph called chord diagram by simply connecting the two preimages of a double point x with an edge:

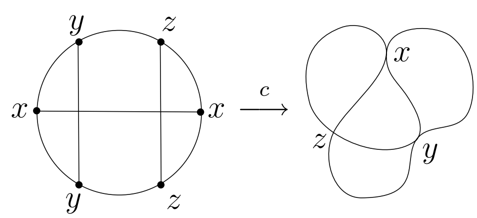

Our aim will be to give a combinatorial characterization of equivalence classes of continuous curves. The precise definitions and many of the constructions involved in the interpretation of a curve as a combinatorial object are founded on deep analytical results such as the Jordan-Schönflies Theorem and the Rotation Principle. We present a brief overview of this background in chapter 1 before turning to the task we want to deal with in chapter 2.

Gauss posed the following problem: Which chord diagrams can be realized by a curve that crosses itself at every double point? Gauss formulated his question in terms of Gauss codes. A Gauss code is nothing but a representation of a chord diagram as word of $2n$ symbols: a Gauss code of the above chord diagram is $xyzxzy$. A lot of research has been done on Gauss' problem and in chapter 2 we are going to present three answers to his question, which are due to Lovász-Marx [6], Rosenstiehl [9] and De Fraysseix-Ossona de Mendez [3]. The first two are classic theorems given in the 1970s: Lovász-Marx define a substructure relation on Gauss codes and give a set of obstructions under this relation that characterizes cross-realizability. Rosenstiehl shows that a chord diagram is cross-realizable iff there exists a partition of the set of chords into two classes that has certain properties. His proof is based on the theory of maps.

The last one is fairly recent (1997) and its elegance has renewed the interest in the subject. De Fraysseix-Ossona de Mendez define an operation called "switch duplication" on crossing curves that preserves cross-realizability. Using this operation they can transform a crossing curve into a touching curve and touching curves are easy to characterize. This idea provides them not only with an elegant original characterization of cross-realizability, but it also turns out that the condition of the Rosenstiehl Criterion is invariant under "switch duplication" and thus they obtain a short proof of Rosenstiehl's theorem.

It turns out, that all three theorems are closely related to the following fundamental observation: On the one hand a curve c with a crossing double point x can be transformed into a curve c' in which x is touching, while on the other hand c can be transformed into a pair of curves c_1 and c_2 that touch at x.

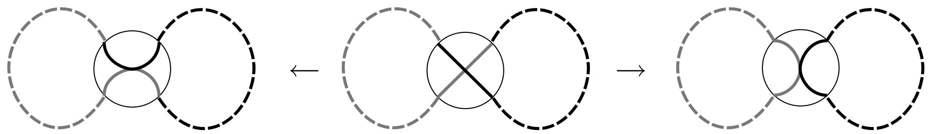

This observation leads us to consider curves that can have both touching and/or crossing double points. To cope with this more general kind of curve we introduce augmented chord diagrams. We use the left-hand side of the above figure to define a switch operation on these objects under which realizability is invariant to obtain a de Fraysseix-Ossona de Mendez-type characterization of this more general kind of curve. We are then able to formulate and prove a Rosenstiehl-type criterion for the realizability of augmented chord diagrams, generalizing the approach of de Fraysseix-Ossona de Mendez. From these generalizations, the theorems by Rosenstiehl and de Fraysseix-Ossona de Mendez follow as a corollary and, following Aigner [1], we give a proof of the Lovász-Marx Criterion from the original Rosenstiehl Criterion. We conclude chapter 2 by showing that for every realizable augmented chord diagram there is exactly one (equivalence class of) realizing curves, and by providing an outlook on how the right-hand side of the above figure might be incorporated into this framework.

Note that we do not consider the application of the "switch approach" to graph drawings instead of curves (see [4]) nor do we consider curves on other surfaces than the plane/sphere (see [1]).

A thrackle is a drawing of a graph G in the plane such that every two edges e_1, e_2 have exactly one point in common: either a common end-point or a common interior point at which they cross.

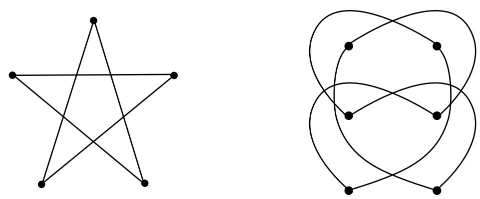

Conway introduced these objects and conjectured in the 1960s that if a graph G has a thrackle drawing then $|E(G)| \leqslant |V(G)|$. This conjecture still remains open and the best known bound to date is $|E(G)| \leqslant \frac{3}{2}(|V(G)| - 1)$ [2]. Early on, Woodall showed [10] that the thrackle conjecture is equivalent to the statement that certain graphs that consist of two circles that have one vertex in common cannot be thrackled. We give a detailed account of Woodall's reduction of the thrackle conjecture. In section 3.4 we turn the Thrackle Conjecture into a question about closed curves in the plane, to which the tools developed in chapter 2 can be applied. We demonstrate this by showing that there are precisely 3 thrackle drawings of C_6 and close chapter 3 by introducing "touching thrackles" and showing that the bound on the number of edges holds in this simpler scenario.

The chapters are more or less independent from one another.

I thank Prof. Dr. Martin Aigner for his support, his encouragement and the interesting further questions to consider he always comes up with. I thank Dr. Mark de Longueville for his time, his enthusiasm and many helpful conversations about the subject. I thank Daria Schymura for several remarks on the structure of the text. Finally, I thank Doris Breuer for her help with the English language.

This document was typeset with TeXmacs (`http://www.texmacs.org`). The figures were drawn with Inkscape (`http://www.inkscape.org`) and then labled with TeXmacs. I would like to thank the authors of both of these free software projects for their work.

Chapter 1
From the Continuous to the Discrete

1.1 Curves and Chord Diagrams

> *The objects we want to study are continuous closed curves in the plane with finitely many double points and no other multi points.*

Two properties of such curves are of particular interest to us:

> *In what order do we encounter the double points when traversing the curve?*
>
> *Does the curve "cross" itself or does it "touch" itself at a given double point?*

Before making the above precise, let us take a look at some examples. The drawings in the top row of Figure 1.1 show curves of the type we want to consider, while the curves shown in the bottom row are not of the type we want to consider. Curves a), b) and c) have three double points each and no other multi points. Curve d) has a single multi point through which it passes three times. Curve e) passes through each point in the interval in the middle twice, and hence has infinitely many double points.

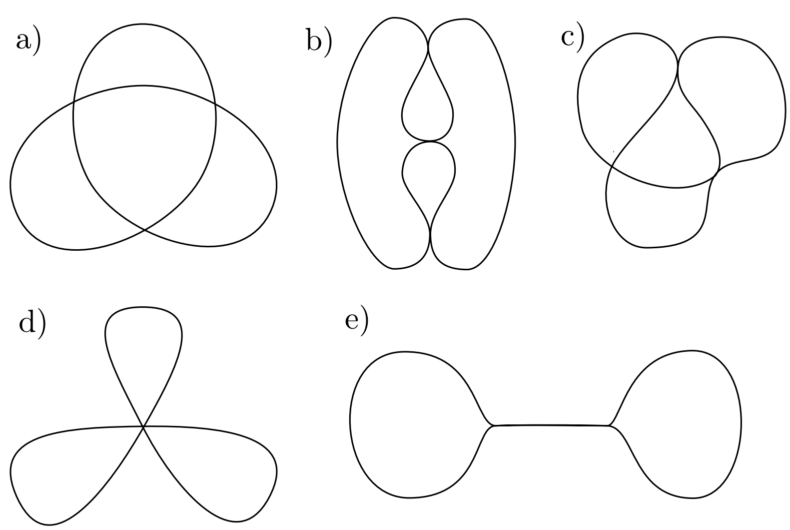

Figure 1.1. We cannot consider a curve as a point set since we need to know how that point set is traversed. Hence only the images of curves are shown. a), b) and c) are of the type we consider, while d) and e) are not.

One way to define this type of curve formally would be the following: A **closed curve** is a continuous function $c\colon S^1 \to \mathbb{R}^2$ from the circle $S^1 = \{x \in \mathbb{R}^2\colon \|x\| = 1\}$ into the plane $\mathbb{R}^2$. A **multi point** of c is an element $x \in \mathbb{R}^2$ with $|c^{-1}(x)| \geqslant 2$. A **double point** of c is a multi point x with $|c^{-1}(x)| = 2$. As already mentioned, we restrict our attention to closed curves that have a finite number of double points and no other multi points: whenever the term "curve" is used in the text these restrictions apply, unless something else is explicitly stated in the text.

However, the above is not the definition we are going to use. Rather, we are going to adopt a definition based on graph drawings which is presented in section 1.2. The reasons for this choice are subtle and will become clear later on. However, the familiar definition above provides the right intuition and is, in fact, synonymous with the definition in section 1.2 unless equivalence classes of curves are considered.

Chord Diagrams and Interlacement Graphs

We will now answer our two questions for each of the valid examples from Figure 1.1. If we traverse a curve once in its entirety, we will, by definition, encounter each double point exactly twice: Every point $x \in c(S^1) \subset \mathbb{R}^2$ corresponds to two points $c^{-1}(x) \subset S^1$ on the circle and all these preimages $\{y \in c^{-1}(x) \colon |c^{-1}(x)| = 2\}$ occur in some "order" on S^1. In Figure 1.2 the double points of the three example curves are labeled. Their preimages on the circle are shown on the left hand side; both with the same label to indicate the double point they correspond to. Additionally, the two preimages of a double point are connected by a chord. The resulting figures are called the chord diagrams of the corresponding curves.

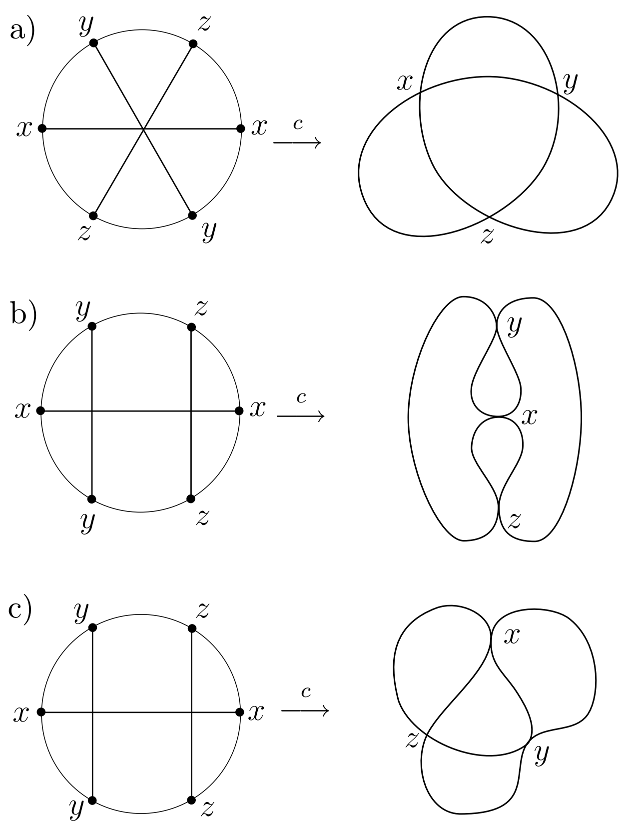

Figure 1.2. The chord diagrams of the three valid examples from Figure 1.1.

It is useful to give here some thought to the notion of "order" we are dealing with, before giving a formal definition of a chord diagram. A **cyclic order** on a finite set S is a circle graph on the vertex set S. A cyclic order is **directed**, if the graph is a directed circle, and it is **undirected** otherwise. A **chord diagram** C is a pair (G_r, G_c) of graphs on a vertex set S, where G_r is an undirected cyclic order (i.e. a circle) and G_c a perfect matching. The edges of G_c are called the **chords** of C, while G_r is called the **rim** of C. Note that the edges of G_r and G_c may coincide, i.e. the disjoint union of both may from a multigraph.

Now, given a curve c the **chord diagram $\mathcal{C}(c)$ of a curve c** is just the cyclic order given by S^1 on the preimages of all double points of c, where the two preimages of a given double point x are matched. Every abstract chord diagram can be represented visually as a circle with straight-line chords, as shown above. It is useful to think about chord diagrams visually.

The chord diagram of a curve c is also called its **Gauss code**. The classical way of defining a Gauss code is this: A **2-word** over the alphabet $\{1, ..., k\}$ is a word of length $2k$, in which every $1 \leqslant i \leqslant k$ occurs exactly twice. We can obtain a 2-word ω over $\{1, \ldots, k\}$ from a curve c with k double points, by

i. labeling all double points with distinct numbers $i \in \{1, \ldots, k\}$,

ii. picking a starting point $x \in S^1$ and an orientation of S^1, and

iii. traversing c once, noting down the label of each double point we pass.

As the result ω depends on the choice of starting point and direction we now consider the equivalence class $[\omega]_\sim$, where $\sim$ is obtained by identifying 2-words that differ up to a cyclic shift and/or reversal. This equivalence class $[\omega]_\sim$ is then called the Gauss code of c. It is clear that chord diagrams and Gauss codes are two different representations of the same concept and we are going to use the two interchangeably.

We call two chords x, y of a chord diagram C **interlaced**, if their corresponding straight lines cross. In terms of Gauss codes, this means that there is a 2-word ω representing C of the form $\omega = x \ldots y \ldots x \ldots y \ldots$. Given a chord diagram C, the **interlacement graph** $\mathcal{I}(C)$ is the graph on the set of chords, in which two chords are adjacent iff they are interlaced. For curves c we will write $\mathcal{I}(c) := \mathcal{I}(\mathcal{C}(c))$. $\mathcal{I}(c)$ is a graph on the set of double points of c.

The interlacement graph of example a) is the complete graph on $\{x, y, z\}$. In the interlacement graph of both b) and c), x is adjacent to both y and z, while y and z are not adjacent.

We call a chord diagram C **realizable** iff there exists some curve c with $\mathcal{C}(c) = C$. If a chord diagram is realized at all by some curve c, it is obviously realized by infinitely many different curves. Of course many of these will be "essentially the same" for our purposes, which is why we need to consider suitable equivalence classes of curves. The definition and study of these equivalence classes will take up much of this chapter. The far more important issue of characterizing all realizable chord diagrams will be dealt with in detail in chapter 2.

Touching and Crossing Double Points

Let us now turn to the second question. When do we consider a double point x to be touching, when do we consider x to be crossing? Well, intuitively, x is touching if c looks as shown on the left-hand side of Figure 1.3 and x is crossing if c looks as on the right-hand side in a small neighborhood of x.

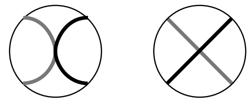

Figure 1.3. Touching and crossing double points.

Let us sketch a formal definition of "touching" and "crossing". Given a double point x, there exists a small neighborhood N around x that contains no double point other than x, because there are only finitely many double points. Also, N can be chosen such that $c^{-1}(N)$ has exactly 2 components. Restricting c to each of these two subsets of S^1 in turn, we obtain two segments c_1 and c_2 of c. In Figure 1.3 the two segments are shown in different colors. Now, c_1 divides N into two connected regions. x is **touching** iff c_2 meets only one of these regions and it is **crossing** iff c_2 meets both of these regions. We call this the **type** of the double point x.

This is a "sketch" of a definition rather than a proper definition, because, to see that the terms are well-defined, we need among countless other analytical details the fact that c_1 divides N into two regions, which implies the Jordan Curve Theorem. So the "obvious" fact that a double point is either touching or crossing turns out to be closely tied to a deep analytical theorem. By and large we will take these "obvious" facts for granted, however we will take a tour of the analytical machinery involved in section 1.3.

For the sake of convenience we agree on the following convention: whenever the image of a curve c in a figure "looks like" the left-hand side of Figure 1.3 in a small neighborhood of a double point, that double point is understood to be touching and the segments of c are understood to be as defined by the colors in Figure 1.3. Accordingly, a double point is understood to be crossing if it "looks like" the right-hand side of Figure 1.3.

Following this convention, the double points of the example curves in Figure 1.2 are as listed in the following table:

	a)	b)	c)
x	crossing	touching	touching
y	crossing	touching	touching
z	crossing	touching	crossing

Equivalence Classes of Curves

As already mentioned, none of the figures uniquely define a curve: There are several curves c realizing a given figure, even when it is understood how each double point is supposed to be traversed (i.e. what the segments in a small neighborhood of a double point are supposed to be). In fact, while there are infinitely many functions $c\colon S^1 \to \mathbb{R}^2$ with k double points, there are only finitely many chord diagrams with k chords and only finitely many vectors in $\{\text{crossing, touching}\}^k$. It stands to reason that we will want to consider equivalence classes of such functions c instead of individual functions. We will now develop some intuitive notions about when we would like to regard two curves c_1 and c_2 as essentially the same, in the context of the types of curves we consider and in the context of the properties we are interested in.

i. *We identify curves c_1 and c_2 that differ only up to reparametrization.* This means that c_1 traverses the same set of points as c_2 but at a different "speed" (cf. Figure 1.4). This obviously leaves the order and the type of the double points invariant.

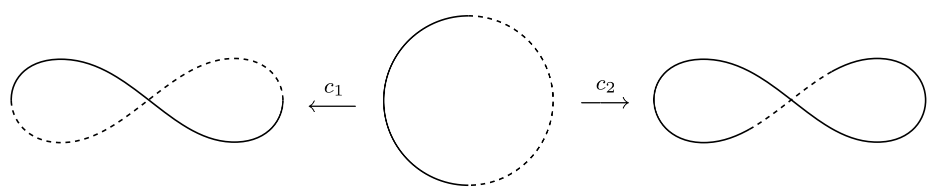

Figure 1.4. The curves c_1 and c_2 the same image at different speeds: the image of the dashed segment of S^1 under c_1 is much longer than the image under c_2.

ii. *We identify curves c_1 and c_2 that differ only up to reversal.* The images of c_1 and c_2 are identical but are traversed in opposite directions. This corresponds to the fact that we regard the rim of a chord diagram as undirected.

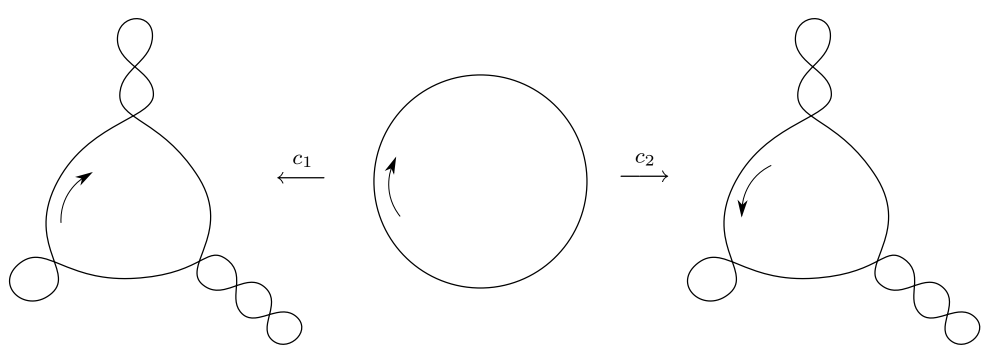

Figure 1.5. c_1 is the reverse of c_2.

iii. *We identify curves c_1 and c_2 that differ only up to a homeomorphism of the plane.* If we can stretch and squeeze, translate and scale, rotate and reflect the entire plane to transform c_1 into c_2, we will identify the two. Note that the reflection of a curve and its reversal are two different operations.

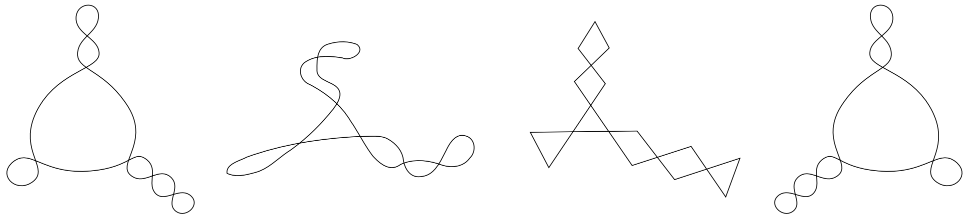

Figure 1.6. For any two of these images there is an automorphism of the plane transforming the one into the other.

iv. *We identify curves c_1 and c_2 that differ only up to polar transformation.* Homeomorphisms do not change the region of the plane that is "outside" of our drawings. However, two curves that differ only with regard to the region that is "outside", i.e. unbounded, have the same chord diagram and the double points have the same type.

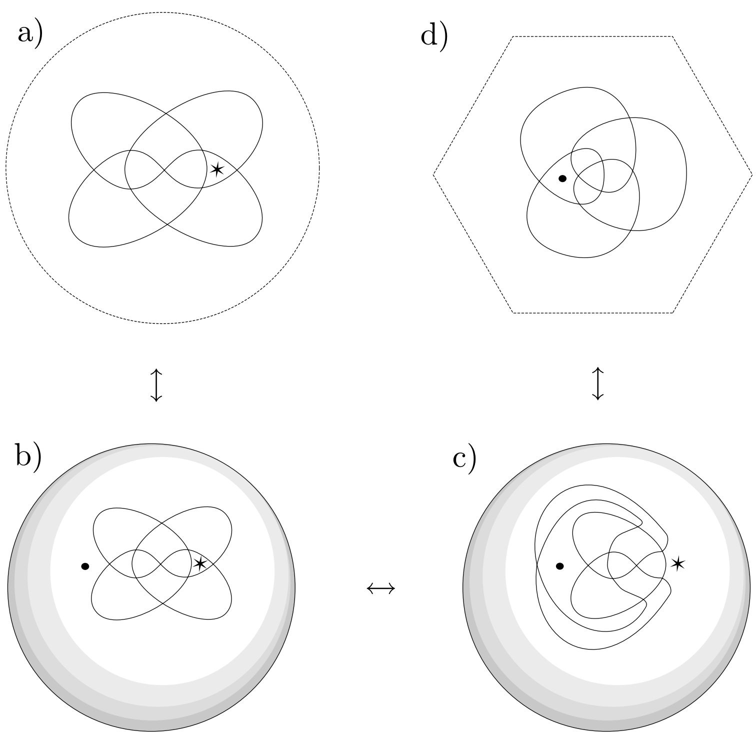

Figure 1.7. Polar transformation. a) We are given a curve in the plane which we want to transform such that the face that contains the star is "outside". We start by picking a large circle that contains the entire curve. By identifying all points on this circle (or by using a more elaborate construction such as the stereographic projection) we obtain a drawing of the curve on the sphere such as the one shown in b). The dot represents the point to which the circle has been shrunk. c) We can now "pull" segments of the curve around the back of the sphere, i.e. apply an automorphism of the sphere. The crucial step is that we can remove the point marked with the star from the sphere to obtain a topological space homeomorphic to the disk shown in d).

We have now identified several operations on curves that appear to have no effect on the two properties of curves we are interested in. Based on these observations we could now give the following definition, if we were to stick to the view of a curve as a continuous function. Let i: $\mathbb{R}^2 \to S^2$ be the (inverse of the) stereographic projection. Two curves c_1 and c_2 are equivalent, iff there exist homeomorphisms h_1: $S^1 \to S^1$ and h_2: $S^2 \to S^2$ such that

$$h_2 \circ i \circ c_1 \circ h_1 = i \circ c_2$$

This means, we first regard c_1 and c_2 as curves $i \circ c_1$ and $i \circ c_2$ on the sphere and then allow these curves to differ up to homeomorphisms of the sphere by introducing h_2. The combined effect is that we make identifications iii. and iv. as well as i. However, ii. is not included, as the example given in ii. shows, so we have to introduce h_1 to allow for the reversal of curves. As already mentioned we will instead adopt the definition presented in the next section, though.

We will often identify a curve with its equivalence class. In most cases the term "curve" will refer to an equivalence class of curves. It will be explicitly mentioned when we make an exception to this rule or when we want to stress that we are talking about equivalence classes.

1.2 Curves and Graphs

Our definitions of a curve and what it means for a double point to be touching or crossing were analytic. The fact that we consider only curves with a finite number of double points and no other multi points allows us to take also a discrete point of view. The idea is that the image of a curve c defines the drawing of a graph $\mathcal{G}(c)$ in the plane, the vertices of which are the double points, while the way the curve traverses its image defines an Euler tour through that graph.

For this to be correct we need to carefully define what exactly an Euler tour is. Normally one would define an Euler tour to be a walk that contains every edge exactly once. However, this definition is too coarse for our purposes here. Consider Figures 1.8a) and 1.8b). The graph defined by the image of these two curves is in both cases a single vertex x with two loops a and b (Figure 1.8c). If we consider a tour as an alternating sequence of vertices and edges (with the usual properties) then the Euler tour is $xaxbx$ in both cases, even though the curves are different: the curve on the left traverses one edge in the opposite direction than the curve on the right. To capture this difference in our concept of an Euler tour, we will consider an edge $e = vw$ as a pair of **half-edges** h_1 and h_2. The order of the two half edges in the tour then gives the direction in which we traverse the corresponding edge. In our example, we consider the edge a as a pair of half-edges a_1 and a_2 and the edge b as pair of half-edges b_1 and b_2 as shown in Figure 1.8d). The curve a) then corresponds to the Euler tour $a_1\, a_2\, b_2\, b_1$ and the curve b) corresponds to the Euler tour $a_1\, a_2\, b_1\, b_2$ – so we can distinguish between the two.

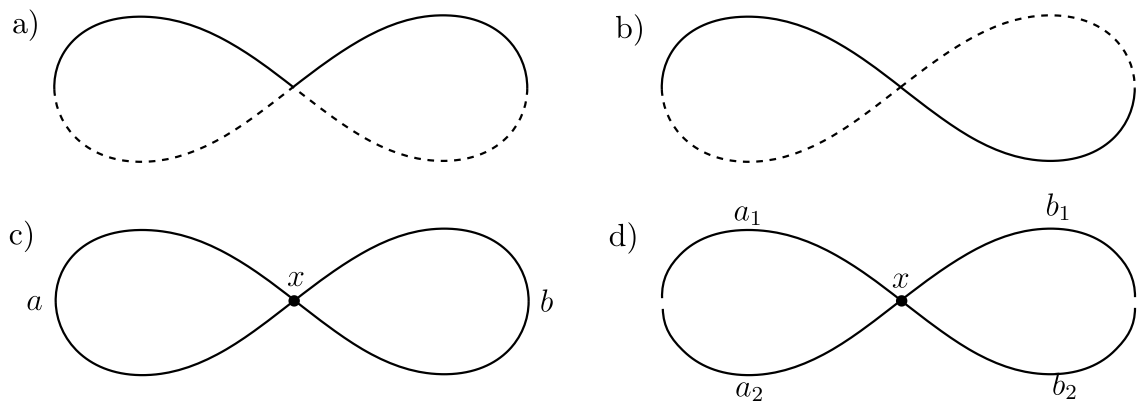

Figure 1.8. The curves a) and b) are different but they define the same plane graph c). The Euler tours of a) and b) in c) are both $x\,a\,x\,b\,x$. Considering every edge as a pair of half edges as in d) we can distinguish between the tours $a_1\, a_2\, b_2\, b_1$ and $a_1\, a_2\, b_1\, b_2$.

We are now going to give precise definitions of all the concepts involved and mention some related theorems and lemmata. No attempt is made to prove the presented results, however, as a sound derivation of the related theory is out of the scope of this work.

Multigraphs

A **multigraph** G is a triple (V, H, E) of a set H, the elements of which are called **half-edges** and two partitions V and E of this set H. V is an arbitrary partition of H into n classes which are called **vertices**. E is a partition of H into classes of size 2 that are called **edges**. The pair (V, E) forms a multigraph as usual, where $v \in V$ is incident to $e \in E$ iff there exists a half-edge $h \in v \cap e$. We say a half-edge h is incident with a vertex v if $h \in v$ and h belongs to an edge e if $h \in e$. If for two half edges h and h' we have $\{h, h'\} \subset v$ for some vertex v, we say that h and h' share v. Similarly h and h' share an edge e if $\{h, h'\} \subset e$. Note that two half-edges can share at most one vertex and at most one edge. However they can share both a vertex and an edge at the same time, in which case they form a loop.

Given a fixed multigraph G a **walk** is a sequence $w = h_1 \ldots h_{2k}$ of half-edges with the property that h_i and h_{i+1} share an edge if i is odd, and that they share a vertex if i is even. The shared vertices and edges are said to be contained in w and it is this alternating sequence of vertices and edges that is usually called a walk. A walk is **closed** if h_1 and h_{2k} share a vertex. We identify walks that differ only up to cyclic shifts, i.e. we identify w with all walks of the form $h_{2j+1} \ldots h_{2k} h_1 \ldots h_{2j}$. We call a walk **undirected** if we also identify walks that differ up to reversal; otherwise we call the walk **directed**. By default a walk is considered to be undirected. A walk is **simple** if it contains no edge more than once. Note that a simple walk can be seen as a cyclic order on a set of half-edges. A **cycle** is walk that contains no vertex more than once. A simple, undirected walk that contains all vertices and all edges of a multigraph is called an **Euler tour**. If a multigraph G has an Euler tour, it is called **Eulerian**.

Drawings of Multigraphs

An **arc** is either an open arc or a closed arc and these are defined as follows: an **open arc** is the image (!) of some continuous, injective function $f\colon [0,1] \to \mathbb{R}^2$. The points $f(0)$ and $f(1)$ are the **endpoints** of the arc and the image of the open interval $(0,1)$ under f is the **interior** of the arc. Note that the endpoints and the interior of an arc are independent of the choice of f. An open arc is said to **connect** its endpoints. A **closed arc** is the image of a continuous, injective function $f\colon S^1 \to \mathbb{R}^2$ with a given endpoint $x \in f(S^1)$. The interior of an open arc is $f(S^1) \setminus x$. A set $X \subset \mathbb{R}^2$ is **connected** if for any two points $x, y \in X$ there exists an arc connecting the two. A **segment** of an arc given by some function f is the image of the restriction of f to some connected subset of the domain of f, i.e. an "interval".

Intuitively, a plane drawing of a multigraph G is a drawing of G in the plane such that no two edges meet, except at a common vertex. We will now give a precise definition. To keep things simple we consider only (drawings of) multigraphs G that have no isolated vertices. In particular we do not consider the graph that consists of a single vertex and no edge, but we do allow the graph that consists of a single loop. A **drawing** or **embedding** of a multigraph G in the plane is a function d mapping each half-edge h of G to a simple open arc $d(h)$ in the plane, with the following properties: a) if $h \neq h'$, then $d(h)$ and $d(h')$ do not have an interior point in common, b) for every edge e of G there is a point $d(e)$ that exactly the two half-edges belonging to e have in common, c) for every vertex v of G there is a point $d(v)$ which exactly the half-edges incident with v have in common and d) the points $d(a)$ for $a \in V \cup E$ are all distinct. Note that a consequence of this definition is that the $d(e)$ and the $d(v)$ are endpoints of the corresponding (sets of) half-edges. A pair (G, d) of a multigraph G and a plane drawing d of G is a **plane multigraph**. If G has a plane drawing, it is called **planar**. We will sometimes denote a plane multigraph by a single letter G, without assigning a separate symbol to the actual drawing d.

Given a plane multigraph (G, d), we can consider the subset $\bigcup_{h \in H} d(h)$ of the plane. In abuse of notation we will denote this set with $d(G)$. Similarly, we can consider an edge $e = \{h, h'\}$ of G as a subset $d(h) \cup d(h')$ of the plane, which turns out to be an open or closed arc. The vertices v, in turn, are the points $d(v) \in \mathbb{R}^2$. The **faces** of (G, d) are the connected components of $\mathbb{R}^2 \setminus d(G)$. Note that the faces of a multigraph are open. The terms **closure**, **interior**, and **boundary** have the usual topological definition. However, the following lemma allows us to attach some combinatorial meaning to these topological terms.

1.1. Let G be a plane multigraph.

i. The interior of an edge e is either contained in or disjoint from the boundary of a face f.

ii. An edge e is in the boundary of $1 \leqslant k \leqslant 2$ faces. $k = 1$ iff e is a bridge.

iii. The boundary of a face can be traversed by a closed walk w that contains every edge at most twice. For an edge e contained in w, we have: e is traversed twice iff e is a bridge. If e is traversed twice by w, it is traversed in opposite directions.

These statements are fairly intuitive. However, as with most results presented in this chapter, we are not going to give a proof of 1.1.

A Curve as a Graph Drawing with an Euler Tour

We can now define what a curve is supposed to be. A **curve** c is a triple (G, d, τ) where G is a 4-regular multigraph, d a plane drawing of G and τ an Euler tour of G. The **double points** of c are the vertices of G. If c is a curve we denote the graph G with $\mathcal{G}(c)$.

This really is our *definition* of a curve. The analytical concept mentioned in the previous section does not serve as a definition, even though a continuous function $c: S^1 \to \mathbb{R}^2$ with the given properties certainly does define a such triple (G, d, τ). The reason for this choice will become clear later on, when we count the number of different realizations of a given chord diagram.

If we consider a curve as a triple (G, d, τ), what, then, is a chord diagram? The chord diagram of a curve is uniquely determined by the pair (G, τ), i.e. it is not dependent on the drawing d, as it can be constructed in the following fashion: We duplicate every vertex of G and insert a chord between the two copies of each vertex. If $h_1, \ldots, h_4$ are the half-edges incident with a vertex v and the Euler tour is of the form $\tau = h_1 h_2 \ldots h_3 h_4 \ldots$, then one copy of v is incident with the half-edges h_1 and h_2 while the other is incident with h_3 and h_4.

Conversely, given a chord diagram C we can obtain the graph G by contracting every chord of C and we can obtain τ by deleting every chord of C (and considering τ as the cyclic order on the half-edges given by the resulting circle). All in all, we can view the combinatorial part of a curve either as a pair (G, τ) or as chord diagram C and both of these views are equivalent.

In this setting, let us again consider what it means for two chords x and y to be interlaced. The four vertices incident to x and y divide the rim of C into four segments. x and y are interlaced iff all four segments connect vertices of different chords. x and y are not interlaced iff only two segments connect vertices of different chords (see Figure 1.9). Consider the graph G' obtained from C by contracting x and y and deleting all other chords. x and y are interlaced iff there are 4 edge-disjoint walks from x to y in G' and they are not interlaced iff there only 2 edge-disjoint walks from x to y. Note that the above observation about edge-disjoint walks does not hold for G as Figure 1.9e) shows.

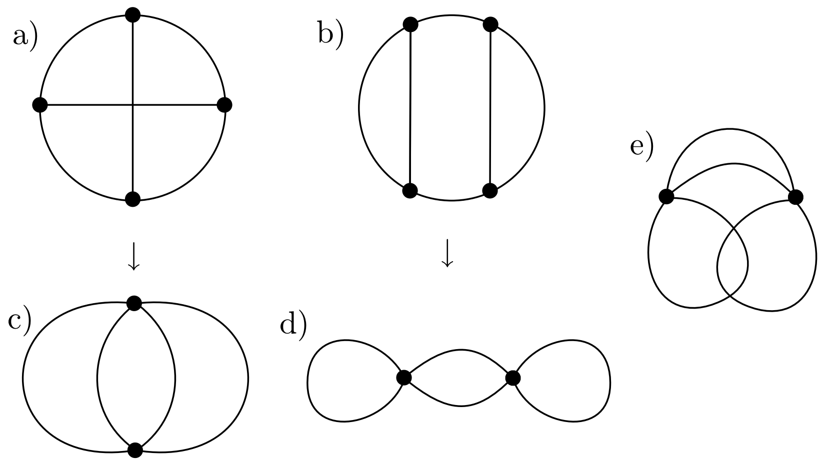

Figure 1.9. In chord diagram a) the two chords are interlaced. Contracting both one obtains the multigraph c) in which there are four edge-disjoint walks connecting the two. In chord diagram b) the two chords are not interlaced and contracting both we obtain multigraph d) in which there are only two edge-disjoint walks between the two. e) Shows two non-interlaced double points of a curve that are nonetheless connected by four edge-disjoint walks in the induced multigraph G.

Equivalence Classes of Graph Drawings

As indicated in the previous section we want to consider equivalence classes of curves. To that end we need to develop a concept of equivalence for multigraph drawings (G, d). One objective is to identify multigraph drawings that differ only up to polar transformation. For this purpose we regard a drawing d of a multigraph G in the plane as a drawing $i \circ d$ of G on the sphere, where $i\colon \mathbb{R}^2 \to S^2 \setminus \{(0, 0, 1)\}$ is a fixed homeomorphism such as the (the inverse of) the stereographic projection. The definition of a drawing on the sphere is just the definition of a plane drawing where the plane $\mathbb{R}^2$ has been replaced with the sphere S^2. The results about faces and boundaries hold just the same. Note that conversely, given a drawing d of G on the sphere, we can always find a point $x \notin d(G)$ and construct a homeomorphism $j\colon S^2 \setminus \{x\} \to \mathbb{R}^2$ using x as north pole, such that $j \circ d$ gives us a drawing of G in the plane.

Let d and d' denote two drawings of a multigraph G on the sphere. d and d' are **equivalent** iff there exists a homeomorphism $\varphi\colon S^2 \to S^2$ such that $\varphi \circ d = d'$. This means that for any half-edge h we have $\varphi(d(h)) = d'(h)$, which implies that φ maps edges onto edges and vertices onto vertices. This definition has the desired effects outlined in section 1.1.

Note, however, that this definition does not take possible symmetries of G into account, so that we do not have to deal with any non-trivial automorphisms G may have. Figure 1.10 shows two non-equivalent drawings of the same graph. In 1.10a) there is a face that has both the loop incident with vertex y and the loop incident with vertex z on its boundary. In 1.10b) there is no such face.

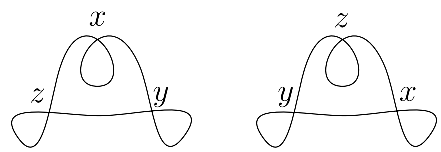

Figure 1.10. Two non-equivalent drawings of a multigraph with three vertices x, y, z.

Having now defined a concept of equivalence for graph drawings, we can finally take to the task of defining a concept of equivalence for curves. It does not suffice for our purposes to require that two curves (G, d, τ) and (G, d', τ') have equivalent plane graphs (G, d) and (G, d') as the example from Figure 1.8 shows: We regard $(G, d, a_1\, a_2\, b_1\, b_2)$ and $(G, d, a_1\, a_2\, b_2\, b_1)$ as different even though the graph drawings are not only equivalent but identical. So we define two curves (G, d, τ) and (G, d', τ') to be equivalent iff (G, d) and (G, d') are equivalent and $\tau = \tau'$. Note that we defined Euler tours τ as undirected walks, so these are equivalence classes of curves that differ up to cyclic shifts and/or reversal.

Recall that the chord diagram of a curve depends only upon G and τ, so equivalent curves have the same chord diagram by definition. It is far from clear, though, whether the *type* of a given double point x is the same for all curves in a given equivalence class. However, our definition of equivalence is such that it makes sense to speak of "the double point x of a given equivalence class of curves". Had we defined a curve as a function this would not be possible as a double point x in c_1 might have several counterparts in an equivalent curve c_2 and thus we would have to take automorphisms of the underlying graph $\mathcal{G}(c_1)$ into account.

Two graph drawings (and consequently two curves) are equivalent if there exist certain homeomorphisms. If we want to show that such homeomorphisms do not exist, we can look at the effect the composition of a graph drawing d with a homeomorphism has on the edges and faces of d in the hope of developing necessary criteria for the equivalence of graph drawings. The following results about the effect of a homeomorphism on the faces and boundaries of a graph drawing are immediate and of combinatorial interest:

1.2. Let G be a plane multigraph and $\varphi: \mathbb{R}^2 \to \mathbb{R}^2$ a homeomorphism.

i. An edge e is in the boundary of a face f of G iff $\varphi(e)$ is in the boundary of $\varphi(f)$ in $\varphi(G)$.

ii. A walk w is the boundary-walk of a face f of G iff $\varphi(w)$ is the boundary-walk of $\varphi(f)$ in $\varphi(G)$.

Analogous statements hold for multigraphs embedded on the sphere or (any other surface).

Obviously 1.2 is of great combinatorial interest as it states that the combinatorial relationship between edges and faces is preserved by homeomorphisms. It is thus well-defined to speak of "the boundary-walks" of an equivalence class of curves. 1.2 gives rise to the hope that a combinatorial characterization of the equivalence of graph drawings can be found, and indeed this is the case. However the path to such a characterization is far too long to be presented here. In the next section, we will merely take a look at some important milestones along the way.

1.3 "Obvious" Analytical Theorems

One of the fundamental results on plane graphs is Euler's formula.

1.3. Euler's Formula

If G is a plane graph with v vertices, e edges and f faces, then

$$v - e + f = 2$$

Considering a simple closed curve as a plane drawing of the graph with a single vertex and one loop, the Jordan Curve Theorem becomes a special case of Euler's formula.

1.4. Jordan Curve Theorem

A simple closed curve $c\colon S^1 \to \mathbb{R}^2$ divides the plane into two connected components and c is the boundary of both. One component F is bounded (in the sense of $\exists r \forall x \in F\colon \|x\| < r$) while the other is unbounded.

The Jordan Curve Theorem appears to be so "obvious" that one feels inclined to take it for granted. In fact, Euler took it for granted when formulating his theorem: Euler died in 1783, more than 50 years before Jordan was born in 1838 and more than a century before the first correct proof of the Jordan Curve Theorem was given by Veblen in 1905 [8].

As was already indicated, most of the results in this thesis are based on analytical facts like the Jordan Curve Theorem. The most crucial of these is that for every plane graph there exists a set of neighborhoods around the vertices and edges that "look as we would expect" (see 1.9). Unfortunately these "obvious" facts are too deep for us to prove within the scope of this work. So we will confine ourselves to presenting some essential theorems, not all of which are obvious. The most remarkable of these is the Heffter-Edmonds-Ringel Rotation Principle 1.10.

This section is based on the book by Mohar and Thomassen [8], to which we refer the interested reader for details.

Constructing Homeomorphisms

Showing that two graph drawings are equivalent requires the construction of a homeomorphism. The premiere tool for the construction of homeomorphisms is the Jordan-Schönflies Theorem.

1.5. Jordan-Schönflies Theorem

Let $c\colon S^1 \to \mathbb{R}^2$ and $c'\colon S^1 \to \mathbb{R}^2$ denote simple closed curves in the plane. Then $c' \circ c^{-1}\colon c(S^1) \to c'(S^1)$ is a homeomorphism which can be extended to a homeomorphism $h\colon \mathbb{R}^2 \to \mathbb{R}^2$ of the entire plane.

This in effect says that any two curves without double points are equivalent. The Jordan-Schönflies Theorem is much stronger than the Jordan Curve Theorem as it guarantees the existence of an homeomorphism of the entire plane and the latter follows easily from the former: The identity id: $S^1 \to \mathbb{R}^2$ is a simple closed curve in the plane. The complement of S^1 is the union of the sets $\{x\colon \|x\| > 1\}$ and $\{x\colon \|x\| < 1\}$ both of which are obviously connected and have S^1 as the boundary. The two sets are different components because $\|\cdot\|\colon \mathbb{R}^2 \to \mathbb{R}$ is continuous. As faces are mapped to faces and boundaries are mapped to boundaries under homeomorphisms, the Jordan Curve Theorem now follows from the Jordan-Schönflies Theorem.

Of course we want to obtain a result on the equivalence of closed curves with a finite number of double points, or, respectively, on the equivalence of plane graphs. The Jordan-Schönflies Theorem can be applied to prove the following lemma.

1.6. Lemma

Let x, y be two distinct points in the plane and c_1, c_2, c_3 three simple curves connecting the two, whose interiors are mutually disjoint. Then $c_1 \cup c_2 \cup c_3$ has three faces that are bounded by $c_1 \cup c_2$, $c_2 \cup c_3$ and $c_3 \cup c_1$ respectively. Let F denote the face of $c_1 \cup c_2$ not containing c_3. Then any homeomorphism $h\colon \bar{F} \to h(\bar{F}) \subset \mathbb{R}^2$ can be extended to a homeomorphism $h'\colon \mathbb{R}^2 \to \mathbb{R}^2$ that maps c_3 onto any given c_3' connecting $h(x)$ and $h(y)$ whose interior is in $\mathbb{R}^2 \setminus F$.

By iterative application of this lemma and a structure theorem for 2-connected graphs, a version of the Jordan-Schönflies theorem for 2-connected graphs can be obtained. Note that for any two isomorphic plane graphs (G, d) and (G', d') there exists a homeomorphism $h\colon d(G) \to d(G')$.

1.7. Jordan-Schönflies Theorem for 2-Connected Graphs

Let (G, d) be a 2-connected plane graph with faces $(F_i)_{1 \leqslant i \leqslant k}$ and corresponding boundaries $(B_i)_{1 \leqslant i \leqslant k}$. Let F_1 denote the unbounded face (the face on the "outside").

1. F_1 is homeomorphic to an open cylinder $\{x \in \mathbb{R}^2\colon \frac{1}{2} < \|x\| < 1\}$ while the F_i for $i \geqslant 2$ are homeomorphic to an open disk $\{x \in \mathbb{R}^2\colon \|x\| < 1\}$.
2. The B_i are simple cycles of G, such that each edge of G occurs exactly twice in all the cycles taken together.
3. Let (G', d') be another plane graph such that G' is isomorphic to G. Let $(F_i')_{1 \leqslant i \leqslant k}$ denote the faces of d' with corresponding boundaries $(B_i')_{1 \leqslant i \leqslant k}$ such that F_1' is unbounded. Then any homeomorphism $h\colon d(G) \to d(G')$ with $h(B_i') = B_i$ can be extended to a homeomorphism $h\colon \mathbb{R}^2 \to \mathbb{R}^2$.

Now, this is already a big step in the right direction. It does not characterize our notion of equivalence for graph drawings entirely, though, as it allows us only to construct homeomorphisms between 2-connected plane graphs and not between general multigraphs embedded on the sphere. Nonetheless, this theorem is an important ingredient in most of the theorems involved in proving the Rotation Principle 1.10, in particular 1.9 which is of interest to us in its own right.

"Nice" Neighborhoods

1.8. Let $G = (V, E)$ be a plane graph. There exists a family $(N_v)_{v\in V}$ of disjoint neighborhoods N_v around every vertex v, with the property that for every edge e there is a point x on e with $x \notin N_v$ for all v. Furthermore there is a family $(N_e)_{e\in E}$ of closed neighborhoods N_e around the set $e \setminus \bigcup_{v\in V} N_v$ for every edge e, such that the N_e are mutually disjoint and disjoint from every N_v and all of G is covered by $\bigcup_{v\in V} N_v \cup \bigcup_{e\in E} N_e$.

Consider the example in Figure 1.11. Note that an edge e may have infinitely many points in common with the boundary of an N_v and that e may pass into and out of N_v several times. In general an N_e may even have infinitely many components. Furthermore 1.8 does not tell us much about how the neighborhoods actually look like.

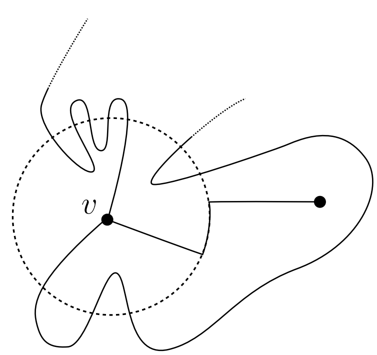

Figure 1.11. The edges of the graph can cross the boundary of N_v multiple times. If this is the case for an edge e, its neighborhood N_e has multiple components.

Using the Jordan-Schönflies Theorem and its analogue for graphs it is possible to strengthen 1.8 considerably.

1.9. Let G, $(N_v)_{v\in V}$ and $(N_e)_{e\in E}$ be as in 1.8.

i. In every N_v there exists a simple closed arc c_v around v that meets every half-edge incident with v exactly once.

ii. The clockwise order of the half-edges on this arc is uniquely determined, i.e. it does not depend on the choice of N_v or c.

iii. Let f be the bounded face of c. There exists a homeomorphism of the plane mapping $c \cup f$ to the wheel with $d(v)$ spokes as depicted in Figure 1.12.

Let h be a half-edge of G and v the vertex h is incident with. Let p_h denote the point where h meets c_v. For every half-edge h let p_h^- be a point on c_v that immediately precedes p_h in the clockwise order and let p_h^+ be a point that succeeds p_h in the clockwise order, i.e. the only element of $\{p_{h'}^-, p_{h'}, p_{h'}^+ \colon h' \text{ a half-edge}\}$ in between p_h^- and p_h^+ is p_h.

iv. For every edge $e = vw$ with half-edges $h_1 = (v, e)$ and $h_2 = (w, e)$ there exist simple closed arcs c_e^1 and c_e^2 connecting $p_{h_1}^-$ with $p_{h_2}^+$ and $p_{h_1}^+$ with $p_{h_2}^-$, respectively. The arcs $\{c_e^1, c_e^2 \colon e \in E\}$ are mutually disjoint, they do not meet any edge of G, and they meet only c_v and c_w in the aforementioned points.

v. We can traverse these arcs in the following way to obtain a simple closed arcs c_e':

$$p_{h_1}^- \xrightarrow{c_e^1} p_{h_2}^+ \xrightarrow{c_w} p_{h_2}^- \xrightarrow{c_e^2} p_{h_1}^+ \xrightarrow{c_v} p_{h_1}^-$$

Here c_v and c_w are traversed in clockwise direction. The bounded face of this simple closed arc is a neighborhood of e. There exists a homeomorphism of the plane mapping this arc and this neighborhood of e to the dumbbell shape depicted in Figure 1.13.

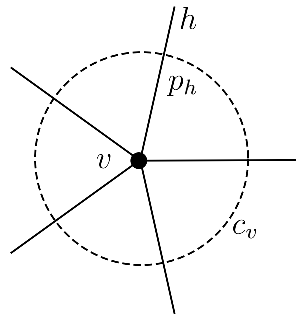

Figure 1.12. The arc c_v is dashed. It meets every half-edge h incident with v in exactly one point p_h.

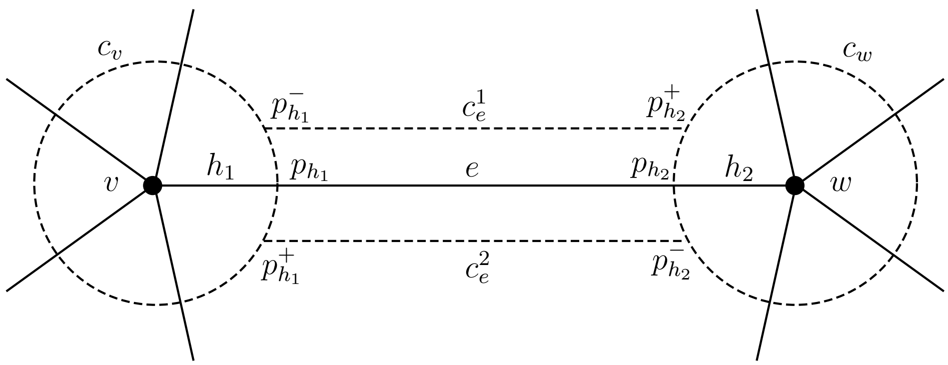

Figure 1.13. The edges in a drawing of a multigraph have a "nice" neighborhood that is homeomorphic to this dumbbell-shape.

In effect 1.9 states that there exist small neighborhoods around the vertices and edges of a plane graph that look as we would expect them to. We are going to use 1.9 quite often, implicitly as well as explicitly, and especially in chapter 3. Note in particular that 1.9.ii says that for an individual curve c the concept of crossing and touching double points is well defined. However, we do not yet know whether the type of a double point x is the same for all curves in a given equivalence class.

Rotation Systems, Surfaces and the Rotation Principle

A **local rotation** π_v at a vertex v of a plane multigraph G is a directed cyclic order of the half-edges incident with v. A **rotation system** is a family $(\pi_v)_{v \in V}$ of local rotations, one for every vertex. We identify rotation systems $(\pi_v)_{v \in V}$ and $(\pi_v^{-1})_{v \in V}$ that differ up to reversal of *all* local rotations. By 1.9 we know that the rotation system of a plane multigraph is a well-defined object.

Let $c = (G, d, \tau)$ be a curve and x one if its double points. We denote the half-edges of x with $h_1, ..., h_4$ such that $\tau = h_1 h_2 ... h_3 h_4 ...$. By 1.9 the drawing d defines a local rotation π_x which is a directed cyclic order of $\{h_1, ..., h_4\}$. The **type** of x is **crossing** iff the pairs $\{h_1, h_2\}$ and $\{h_3, h_4\}$ are interlaced and the type of x is **touching** if they are not interlaced. See Figure 1.14.

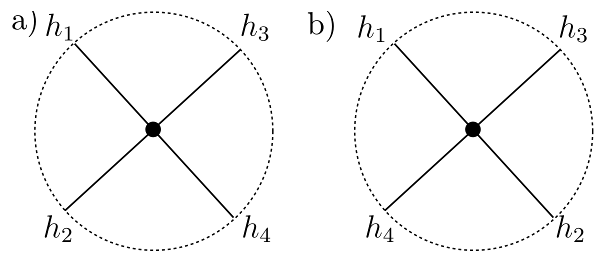

Figure 1.14. If $\tau = h_1 h_2 ... h_3 h_4 ...$ then a) is a touching double point while b) is a crossing double point.

We already defined a chord diagram C to be **realizable**, iff there is a curve c with $\mathcal{C}(c) = C$. The curve c is then called a **realization** of C. We can now refine this concept. C is **cross-realizable** iff there is a realization in which all double points are crossing. C is **touch-realizable** iff there is a realization in which all double points are touching. Let S be the set of chords of C. Let $\vartheta: S \to \{\text{touching, crossing}\}$ be a function assigning a type to each of the chords of C. We say C is **ϑ-realizable** iff there is a realization in which every chord x has type $\vartheta(x)$. We will study the realizability of chord diagrams in chapter 2.

First, we need to see whether it is well defined to speak of the type of a chord in a given equivalence class of curves: Does it make sense to talk about the rotation system of an equivalence class of plane multigraphs? We can observe that the reflection of the plane at an axis swaps the clockwise and anti-clockwise orientation of all simple closed arcs c_v, thus reversing the local rotation at every vertex of a graph G, so our identification of $(\pi_v)_{v \in V}$ with $(\pi_v^{-1})_{v \in V}$ is indeed necessary. But it is not clear whether this identification suffices to make the rotation system of an equivalence class a well defined concept.

It turns out that not only every equivalence class of plane multigraphs has a well defined rotation system, but that the converse is also true: for every rotation system there is an equivalence class of graph drawings on some surface with the given rotation system and, furthermore, the rotation system of a plane multigraph G uniquely defines the equivalence class of G. These astonishing results constitute what is known as the Heffter-Edmonds-Ringel Rotation Principle. To be able to formulate this powerful theorem, we shall now introduce the concept of a surface.

A **surface** is a non-empty, compact, connected, Hausdorff space S with the property that every $x \in S$ has a small neighborhood N that is homeomorphic to the open disk (which in turn is homeomorphic to the plane). Examples of surfaces include the sphere, the torus, the Klein bottle and the projective plane. The plane is not a surface as it is not compact.

We can now consider "drawings of multigraphs without crossings" on arbitrary surfaces. A "multigraph G embedded in the plane" then becomes a "multigraph G **embedded** in a surface S". The definitions of embedding, face, equivalent, etc. carry over directly from the planar case; we simply need to replace $\mathbb{R}^2$ with S. We can interpret the plane multigraphs G we have been considering up to now as embeddings of G on the sphere via the inverse of the stereographic projection – something we needed to do anyway for our concept of equivalence of plane multigraphs.

We call a surface S **orientable** iff it does not contain a Möbius strip. More formally, we say a surface is orientable iff every simple closed arc $a \subset S$ has a small neighborhood $N \supset a$ such that $N \setminus a$ is not connected. We call an embedding of a multigraph G in S **cellular** iff every face of G is homeomorphic to the open disk. Given such a cellular embedding an analogue of 1.9 still holds, the only problem is that we do not know what "clockwise" is supposed to mean on a general surface. However we can find arcs c_v, c_e^1 and c_e^2 as stated in 1.9. Even without a notion of "clockwise" we can define the curve c_e' mentioned in 1.9. Now, we start with a vertex v of G and define an orientation of c_v to be clockwise. Via c_e' this induces a notion of clockwise for all arcs c_w around adjacent vertices w. As we consider only connected multigraphs G, this defines a clockwise orientation around every vertex $v \in V(G)$.

Again, the question arises whether this is well-defined. Not surprisingly, this algorithm yields a well defined result if and only if S is orientable. The condition that S does not contain a Möbius strip excludes problematic situations like the two depicted in Figure 1.15.

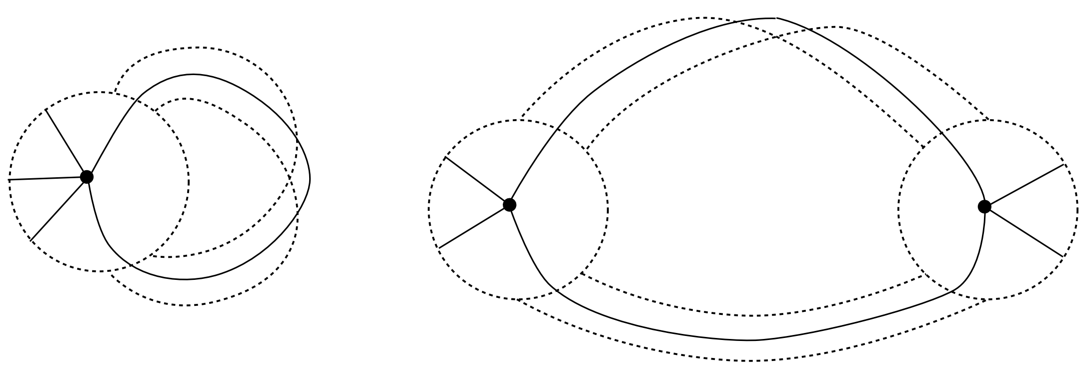

Figure 1.15. If the arcs c_e^1, c_e^2 and points p_h^+, p_h^- from 1.9 are connected as in these two examples, an orientation of the arcs c_v around the vertices cannot be defined consistently. By considering only surfaces that do not contain a Möbius-strip situations such as these can be avoided.

All in all we have an analogue of 1.9 for cellular embeddings of multigraphs on arbitrary orientable surfaces. In particular, the concept of a rotation system is well defined in this case. This is the aspect we were headed for, because now we can formulate the Rotation Principle, which is a beautiful theorem that deserves to be appreciated in its own right.

1.10. Heffter-Edmonds-Ringel Rotation Principle

Let G denote a connected multigraph G containing at least one edge.

i. For every rotation system $(\pi_v)_v$ of G there exists an orientable surface S and a cellular embedding d of G in S such that rotation system of d is $(\pi_v)_v$.

Let $(G,\, d)$ and $(G,\, d')$ be cellular embeddings of G in surfaces S and S' respectively and let $(\pi_v)_v$ and $(\pi'_v)_v$ denote their respective rotation systems.

ii. There exists a homeomorphism $h\colon S \to S'$ such that $d' = h \circ d$ if and only if $(\pi_v)_v = (\pi'_v)_v$.

Recall that according to the identification we made earlier $(\pi_v)_v = (\pi_v^{-1})_v$. The Rotation Principle provides us with a combinatorial characterization of when two multigraph drawings (in the plane or on the sphere) are equivalent.

1.11. Let (G, d) and (G, d') denote plane multigraphs with rotation systems $(\pi_v)_v$ and $(\pi'_v)_v$. (G, d) and (G, d') are equivalent if and only if $(\pi_v)_v = (\pi'_v)_v$.

Let us stress two consequences of 1.11 that are crucial for our purposes: On the one hand 1.11 allows us to count the number of equivalence classes of plane drawings $[(G, d)]$ as we could instead count the pairs $(G, (\pi_v)_v)$ that yield a cellular embedding of G on the sphere – if we knew which pairs $(G,\, (\pi_v)_v)$ have this property. On the other hand we see that all graph drawings in one equivalence class have the same rotation system. As the type of a double point of a curve $(G,\, d,\, \tau)$ was defined in terms of its rotation system, we conclude that it is well defined to speak of the type of a double point x in an equivalence class of curves. Note that the type of the double points in a curve $(G,\, d,\, \tau)$ is dependent on all three parameters, while the chord diagram of a curve is only dependent on G and τ.

1.4 Maps and Dual Graphs

Sometimes, for example in [8], rotation systems are also called maps. These are not to be confused with another important concept that is frequently referred to as a "map": a **map** is a triple $(\pi,\, \tau,\, \sigma)$ of fix-point free involutions on a ground set of **flags**. A flag can be thought of as a "quarter" of an edge; the flags of an edge e correspond to the points $p^-_{h_1}$, $p^+_{h_1}$, $p^-_{h_2}$, $p^+_{h_2}$ in 1.9. The three involutions on the set of flags of course have a geometrical interpretation – as do the rotation system $(\pi_v)_v$ and the partitions V and E on the set of half-edges in our case. Figure 1.16 shows a side-by-side comparison of the two concepts. Note the similarity of Figure 1.16b) to Figure 1.13.

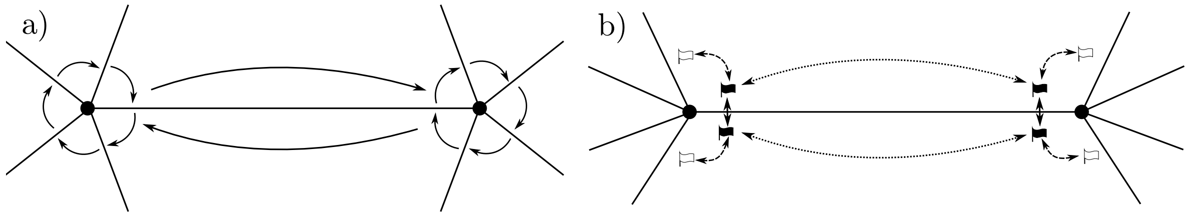

Figure 1.16. a) A local rotation defines a cyclic permutation on the half-edges incident to a vertex. An edge, being a set of two half-edges, can be seen as fixpoint-free involution on the two half-edges. b) In the context of a map, an edge is represented as a set of four flags (shown in black). Three different fixpoint-free involutions are defined on the set of flags, shown with solid, dotted and dashed arrows respectively.

Note that maps are a more general concept than abstract multigraphs with rotation systems, as they allow the representation of cellular embeddings of graphs on non-orientable surfaces. Rosenstiehl's original proof of his theorem 2.14 is based on maps, and indeed these techniques can be applied to characterize Gauss codes of curves on other surfaces than the sphere [1]. We will not consider this approach further as the idea is to pursue a different route based on the work of de Fraysseix and Ossona de Mendez [3].

Let G be a plane multigraph with vertex set V, edge set E and face set F. Its **dual** G^* is defined as follows: the vertices of G^* are the faces of G and the edge set of G^* is the edge set of G. In G^* an edge e is incident to a face $f \in F$ iff e is contained in the boundary of f in G. If e is contained in the boundary of only one face f of G it forms a loop $e = ff$ in G^*. The faces of G^* in turn correspond to the vertices of G, so that we have $G = (G^*)^*$. Note that the edges contained in a boundary-walk of a face f of G are just the edges incident with the vertex f in G^*. A bridge in the boundary-walk of the face f corresponds to a loop incident with the vertex f.

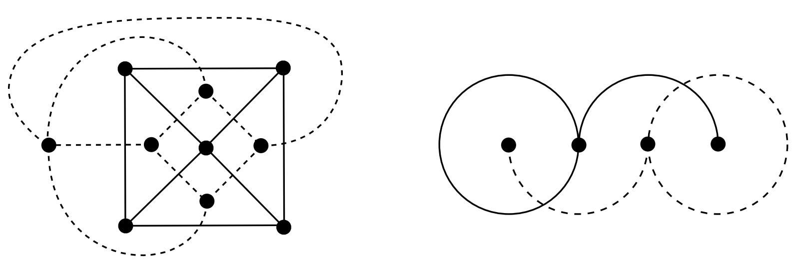

Figure 1.17. Graphs and their dual.

We now want to consider colorings of the faces of G such that no two faces that have a common edge in their boundaries have the same color, i.e. we are interested in colorings of G^* such that no two adjacent vertices of G^* have the same color. We call such colorings proper. We have the following theorem.

1.12. The faces of a curve c can be properly colored with black and white. I.e. if $G = \mathcal{G}(c)$ is the underlying plane graph of c then $\chi(G^*) = 2$.

Proof. We have to show that G^* has no odd cycle. The cycle space of G^* is spanned by the boundary-walks of the faces of G^*. The length of a boundary-walk w in G^* is given by the degree of the corresponding vertex in G. Note that bridges in w correspond to loops in G an are thus counted twice. As every vertex in G has degree 4, every boundary-walk and thus cycle in G^* is even. □

1.5 The Number of Cross-Realizations

As an application of our work on the characterization of equivalence classes of curves, we will now count the number of cross-realizations of a given cross-realizable chord diagram C. Note that we can count these even though we do not yet know which chord diagrams are cross-realizable. The results presented below can also be found in [3].

1.13. Let C be a chord diagram such that the interlacement graph $\mathcal{I}(C)$ is connected. If C is cross-realizable, there is exactly one equivalence class c of curves realizing C.

Proof. Let $(\pi_v)_v$ be the rotation system of a cross-realization of C. Pick any chord x of C. Let the Euler tour defined by C be of the form $\tau = h_1 h_2 ... h_3 h_4 ...$, where h_1, h_2, h_3, h_4 are the half-edges incident with vertex x. As x is crossing, there are two possibilities what the local rotation π_x may be:

$$\begin{aligned} &\text{either} \quad h_1 \xmapsto{\pi_x} h_3 \xmapsto{\pi_x} h_2 \xmapsto{\pi_x} h_4 \\ &\text{or} \quad h_1 \xmapsto{\pi_x} h_4 \xmapsto{\pi_x} h_2 \xmapsto{\pi_x} h_3. \end{aligned}$$

Note that the one alternative is the reverse of the other. Intuitively, these alternatives correspond to whether the segment $h_3 h_4$ crosses the segment $h_1 h_2$ from left to right or from right to left. We assume that the local rotation at x is given by the first alternative and show that then all other local rotations are uniquely determined. Reversal of the local rotation at x has the effect that all other local rotations have to be reversed as well, proving that the equivalence class of $(\pi_v)_v$ is uniquely determined as claimed.

Let y be a chord interlaced with x and denote the half-edges incident with y with h'_1, ..., h'_4 as shown in Figure 1.18. x and y split c into four segments α, β, γ, δ. α and β form a closed curve c' the faces of which can be colored black and white. Without loss of generality we assume that h_4 resides in a white face. The number of times γ crosses c' is uniquely determined by C. As the color of the face the half-edges reside in changes at each crossing, the color of the face h'_3 resides in is uniquely determined.

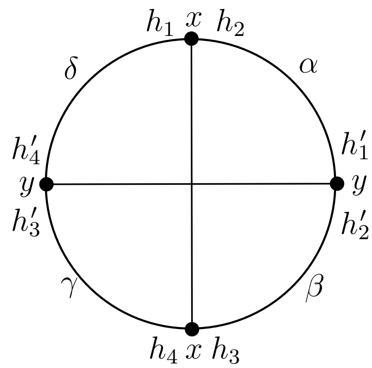

Figure 1.18. Two interlaced chords split a chord diagram and a corresponding curve into four segments α, β, γ, δ. Each contains one half-edge h_i incident with x and one half-edge h'_j incident with y.

On the other hand we can traverse α from h_2 to h'_1. At each of the crossings of c' with itself we pass, the color of the face on the left hand side changes. The color on the left hand side of $h'_1h'_2$ is hence uniquely determined by C and therefore it is uniquely determined on which side of the segment $h'_1h'_2$ the half-edge h'_3 resides. Assuming the plane is oriented clockwise, we put $h'_1 \overset{\pi_y}{\longmapsto} h'_3$, if h'_3 is on the left, and $h'_1 \overset{\pi_y}{\longmapsto} h'_4$ otherwise. □

1.14. Let C be a cross-realizable chord diagram and k the number of components of its interlacement graph Λ. Then

$$\#\text{cross-realizations of } C \;=\; 2^{k-1}$$

Proof. The condition that a double point be crossing determines its local rotation up to reversal. A given local rotation π_x determines the local rotations at all double points in the same component of Λ as x. So at best we can hope to have one degree of freedom per component of Λ which gives us an upper bound of 2^k. The fact that rotation systems $(\pi_v)_v$ and $(\pi_v^{-1})_v$ are equivalent yields the upper bound of 2^{k-1}.

To see that 2^{k-1} different realizations can actually be achieved we employ a simple geometric construction (see Figure 1.19). We use induction on k. Let C be a chord diagram with $k \geqslant 2$ components. There is one component C' that is connected by only 2 edges (in C) to any other components. Split C at these two edges into chord diagrams C' and C'' and reconnect the half-edges to obtain new edges e' and e''. By induction C'' has 2^{k-2} realizations, let c be any one of those. Take a point on e'' and a small neighborhood N around it, insert a realization of C' in which e' is on the boundary of the unbounded face into that neighborhood and reconnect the edges appropriately. Note that C' has only one equivalence class of realizations. Of the representatives of this class for which e' is on the boundary of the unbounded face, we can either pick one in which h_3h_4 crosses h_1h_2 from left to right or from right to left. This choice leads to 2 different realizations as there is another (non-empty) component C'', the local rotations of which are fixed, so we get 2^{k-1} different cross-realizations of C.

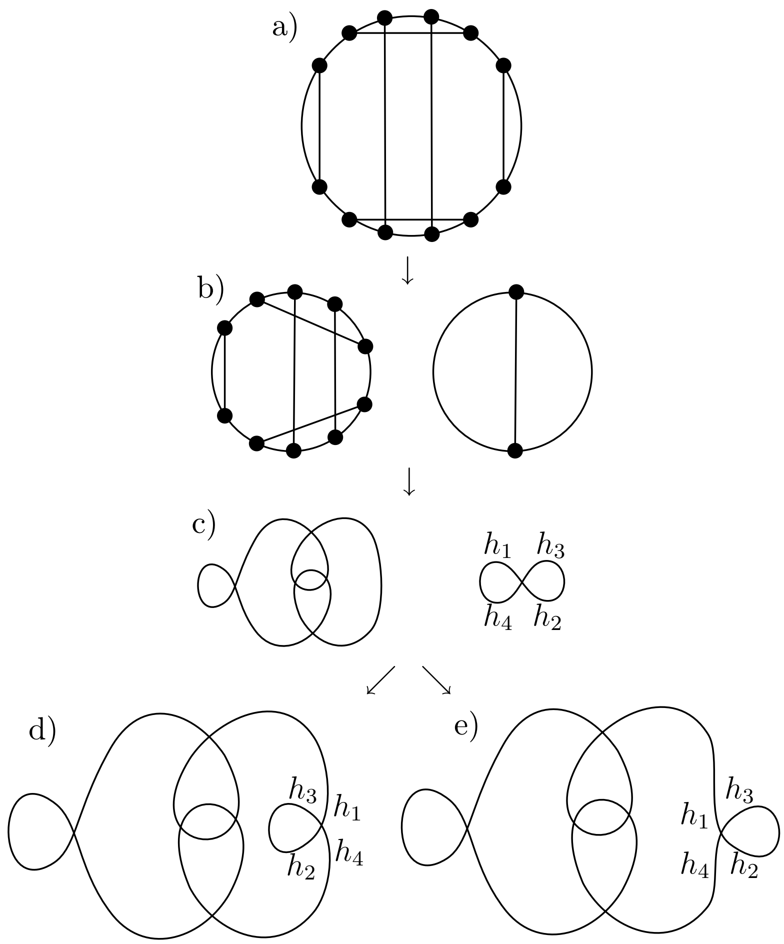

Figure 1.19. Starting with a chord diagram a), we single out a component C' that is connected to other components via only two edges. We split C accordingly and the two resulting chord diagrams are shown in b). One of the two cross-realizations of C'' and the cross-realization of C' is shown in c). These can then be combined in two ways d) and e).

□

Let us take a look at the realizations of the chord diagram $C = xxyyzz$ containing three mutually non-interlaced chords. As $\mathcal{I}(C)$ has three components, there are $2^{3-1} = 4$ realizations which are given in 1.20. Had we chosen the concept of an "equivalence class of curves" as presented in section 1.1, we would have no means to distinguish the last three realizations from one another. In this case, the number of equivalence classes would have been 2 and any general formula for the number of different realizations would have to take symmetries of $\mathcal{G}(c)$ into account.

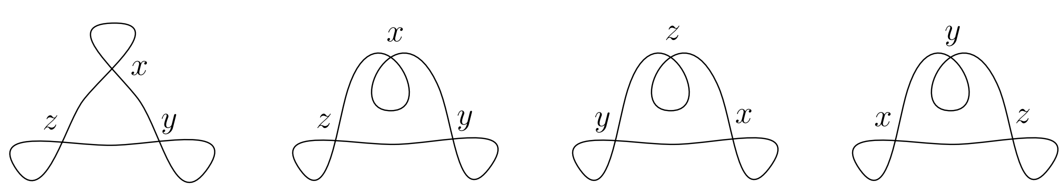

Figure 1.20. The four realizations of the chord diagram $xxyyzz$.

The proof above required a number of rather informal geometric constructions, because we have no combinatorial criteria that can guarantee that choosing a different local rotation when arriving at a new component will not make the multigraph drawing thus defined by its rotation system non-realizable in the plane. We will develop such criteria in the next chapter.

Chapter 2
Characterizations of Curves

2.1 Chord Diagrams and Interlacement Graphs

Closed curves in the plane with finitely many double points and no other multi points are the objects we want to study. As explained in chapter 1 we can associate with each such curve c its chord diagram $C = \mathcal{C}(c)$ and its interlacement graph $\Lambda = \mathcal{I}(C) = \mathcal{I}(c)$ (cf. Figure 2.1).

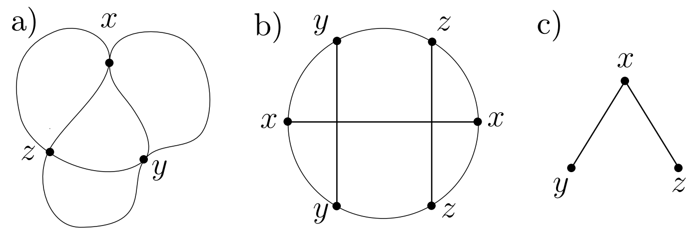

Figure 2.1. A curve a), its chord diagram b) and its interlacement graph c).

A chord diagram is realizable iff it is of the form $\mathcal{C}(c)$ for some curve c. The question we want to answer in this chapter is the following:

Which chord diagrams are realizable?

As we have also seen in chapter 1, we can associate a type with each double point of a curve c: every double point is either touching (such as x and y in Figure 2.1) or crossing (such as z). A chord diagram C is touch-realizable iff it has a realization c that has only touching double points. Similarly, C is cross-realizable iff it has realization c that has only crossing double points. Given a type function ϑ that assigns to each chord x of C a type $\vartheta(x) \in \{\text{touching, crossing}\}$, we call C ϑ-realizable iff there is a realization c in which a double point x has type $\vartheta(x)$. We can now refine our question and ask:

Which chord diagrams are touch-realizable/cross-realizable/ϑ-realizable?

It was Gauss who came up with the concept of a chord diagram in the slightly different form of what is now called a Gauss code (see chapter 1). The original question Gauss asked about these objects was: Which chord diagrams are cross-realizable?

A lot of work has been done on this problem, and we are going to present three different answers to Gauss' question: the theorems by Lovász-Marx, Rosenstiehl and de Fraysseix-Ossona de Mendez. We will follow the approach taken by de Fraysseix-Ossona de Mendez who derive a criterion for cross-realizability from a criterion for touch-realizability. Introducing the concept of augmented chord diagrams, we will then build on their work to give criteria for realizability and ϑ-realizability.

In chapter 1 we interpreted a curve c as a triple $(\mathcal{G}(c), d, \tau)$ of a 4-regular connected multigraph $\mathcal{G}(c)$, a drawing d of $\mathcal{G}(c)$ in the plane and an Euler tour τ of $\mathcal{G}(c)$ and we noted that the pair $(\mathcal{G}(c), \tau)$ uniquely defines the chord diagram $\mathcal{C}(c)$ and vice versa. We are not going to stick too closely to this definition of a curve. Instead we will view a curve as both, a multigraph drawing with an Euler tour on the one hand and as a continuous function on the other, as this makes the text more accessible.

2.2 Matchings at a Double Point

As we know from 1.9 every double point v of a curve c has a "nice" small neighborhood N_v. In essence this neighborhood "looks like" Figure 2.2a). In particular it contains (end-pieces of) the 4 half-edges incident with v. As on any set of 4 elements, there are 3 perfect matchings on this set of half-edges as shown by the edges with box-ends in Figure 2.2b). We will now see that the way c traverses the half-edges of $\mathcal{G}(c)$ gives each of these matchings a special meaning that can be read off from the chord diagram $\mathcal{C}(c)$. For this purpose it is useful to think of a half-edge h as being represented by the point in the plane where h meets the boundary of N_v.

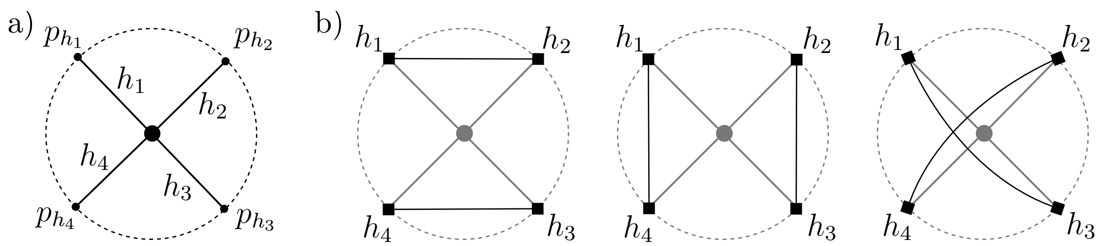

Figure 2.2. Every double point v has a "nice" neighborhood N_v homeomorphic to the one shown in a). b) shows the 3 perfect matchings on set of the four half-edges incident to v.

Note that these are really to be understood as matchings on the abstract set of half-edges incident to a vertex v in the abstract multigraph $\mathcal{G}(c)$. A drawing of $\mathcal{G}(c)$ such as the one given by c defines a local rotation at v. If such a local rotation is given (and the visual representation in Figure 2.2b) does define a local rotation) we can call a matching **touching** iff the matched half-edges are next to each other in the local rotation (as in the first two cases in Figure 2.2b). Otherwise we call the matching **crossing** (the last case in Figure 2.2b). This is the **type** of a matching, and it is given by the drawing d of the multigraph $\mathcal{G}(c)$. We will now seek to classify the matching according to how the graph is traversed, which is given by the Euler tour τ or equivalently by the chord diagram C.

The "nice" neighborhood N_v has the property that $c^{-1}(N_v)$ has two components. We call the two curves given by the restriction of c to these components, respectively, the **local segments** of c at v. Each local segment of v consists of (the end-pieces of) two half-edges incident with v. Thus, the local segments of a curve define a perfect matching on the four half-edges at every vertex v. We call this matching the **local matching** of v. The local segments of the double points in our example from Figure 2.1 are shown in Figure 2.3a).

The two half-edges belonging to one local segment of v are both incident to the same endpoint of the chord v in the chord diagram as Figure 2.3b) shows. For the Euler tour τ this means that half-edges h_1, h_2 are matched by the local matching at v iff they are next to each other in τ and share vertex v.

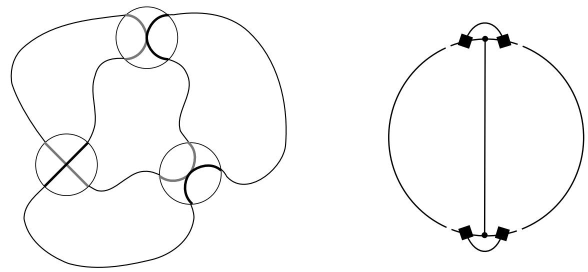

Figure 2.3. a) The two local segments in the neighborhoods of each double point of our example are shown in different colors. b) In general the local matching at a chord v is given by the chord diagram. The two arcs with square ends represent the edges of the local matching.

As already indicated by our example in Figure 2.3a), we will represent the three possible local matchings at a given double point by the symbols shown in Figure 2.4.

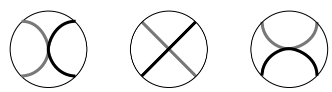

Figure 2.4. We represent the 3 possible local matchings at a given double point by these symbols.

If $c^{-1}(N_v)$ has two components, so has the complement $S^1 \setminus c^{-1}(N_v)$. This time, we call the two curves given by the restriction of c to these components the **global segments** of c at v. The perfect matching on the half-edges incident with v given by the global segments is called the **global matching** of v. Again the global segments and the global matching are given by the chord diagram as shown in Figure 2.5: the global segments are the two components of the rim of the chord diagram after removing the 4 half-edges incident with v. Half-edges h_1, h_2 are matched by the global matching at v iff they are incident to different endpoints of the chord v and connected by a global segment.

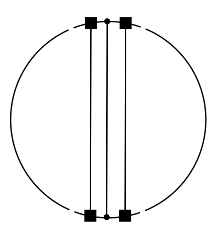

Figure 2.5. The global matching and the global segments are given by the chord diagram.

Figure 2.5 also shows that the global segments can be used to characterize the interlacement interlacement of double points v and w:

2.1. v is interlaced with $w \iff w$ lies on both of the global segments of v

An intuitive description of the role of the local and global matchings is the following: While the local segments are the parts of c inside the neighborhood N_v, the global segments are the parts of c outside of the neighborhood N_v. Correspondingly, we can think of the local matching as giving the pairs of half-edges that are "connected" by c on the inside of N_v and we can think of the global matching as giving the pairs of half-edges that are "connected" on the outside of N_v. Suppose we start at the point p_h on c where the half-edge h meets the boundary of N_v. We can now traverse c in two directions: continuing inside N_v or outside N_v. If we traverse c inside of N_v, we will arrive at the point $p_{h'}$ where the boundary of N_v meets the half-edge h' that is matched with h locally. If we traverse c outside of N_v, we will eventually return to N_v at the point $p_{h''}$ where the boundary of N_v meets the half-edge h'' that is matched with h globally. Correspondingly, we use the symbols shown in Figure 2.6 to represent the 3 possible global matchings at a given double point.

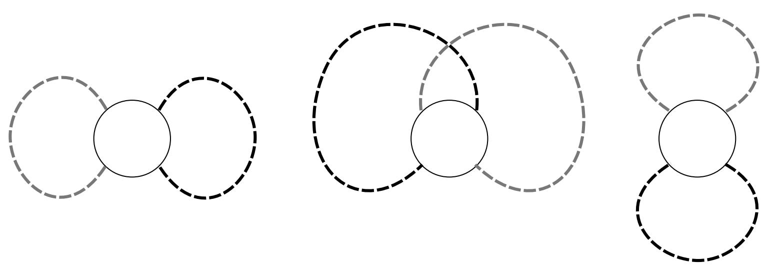

Figure 2.6. We represent the 3 possible global matchings at a given double point by these symbols.

Now, what happens if we change the local matching at a double point v and simply follow the new one in our traversal? To answer this question we consider the effect of substituting each of the possible local matchings at z in our example (Figure 2.7). While 2.7a) and 2.7b) give the image of a single curve, 2.7c) gives the images of two curves. We say 2.7c) is **disconnected**.

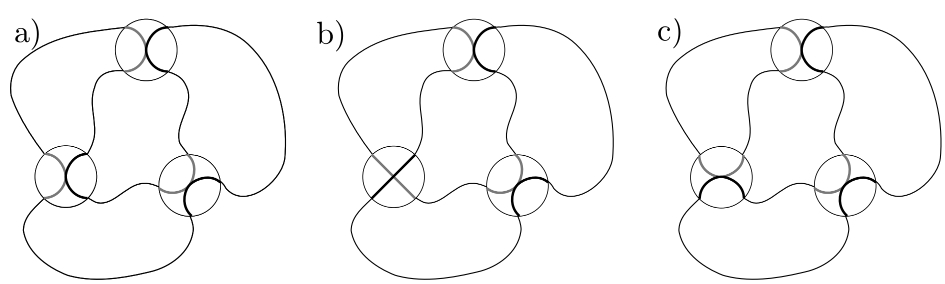

Figure 2.7. The 3 possible local matchings at z give rise to curves a) and b) and the pair of curves c).

The local and global matchings at z in each of the three cases are given in Figure 2.8. We observe that in the case on the right the local and global matchings are the same.

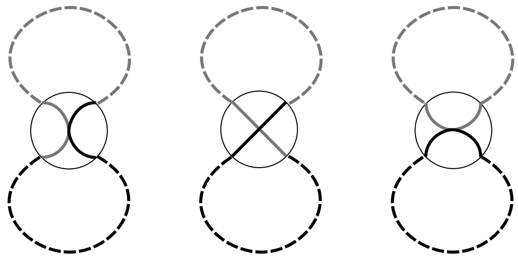

Figure 2.8. The local and global matchings at z in each of the three examples above.

It is immediate that this characterizes the disconnecting local matching in general.

2.2. Let c be a curve and v one of its double points. Let m be a perfect matching on the half-edges incident with v. The result we obtain by substituting m as the local matching at v is disconnected, if and only if m is the global matching of c at v.

It follows that of the three matchings at a double point, one is local and a different one is global. We call the third matching the **diagonal matching**. Of course the diagonal matching is also given by the chord diagram. The local, global and diagonal matchings of a chord diagram C at a chord v are shown side-by-side in Figure 2.9.

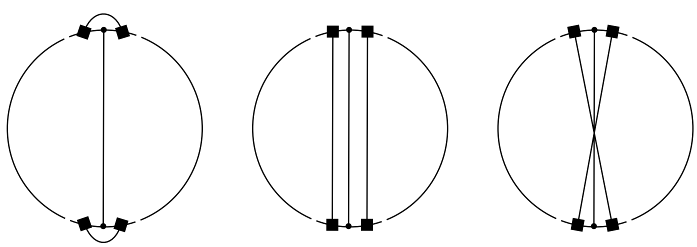

Figure 2.9. The local, global and diagonal matchings at a chord v.

We now have classified the matchings at a double point with regard to two aspects: according to their type in the drawing of c and according to the way they are traversed by c. We now put these two classification in relation to one another. A double point is **locally touching** iff the local matching is touching, it is **locally crossing** iff the local matching is crossing, it is **globally touching** if the global matching is touching, and so on.

2.3 Switches

The explanation given in the last section of what precisely it means to "substitute" or "change" a local matching at a double was intentionally informal. We take a closer look at this operation in section 2.9. For now, we will consider the simpler operation, called "switch", of replacing the local with the diagonal matching, which, by our observations above, does not "disconnect" the curve. This will induce an operation on chord diagrams (resp. interlacement graphs) that preserves realizability.

Intuitively speaking, what is the effect of the switch at v on a curve? We simply take a different turn at v and traverse one of the two global segments in the opposite direction as before (Figure 2.10).

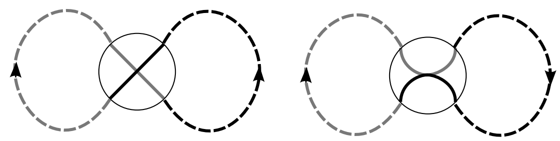

Figure 2.10. The effect of a switch on how a curve is traversed.

The switch at v swaps the local and diagonal matching at v. Thus, the switch at a fixed double point v is an involution. Using our visual notation for local and global matchings Figure 2.11 lists the effect of the switch operation in all cases.

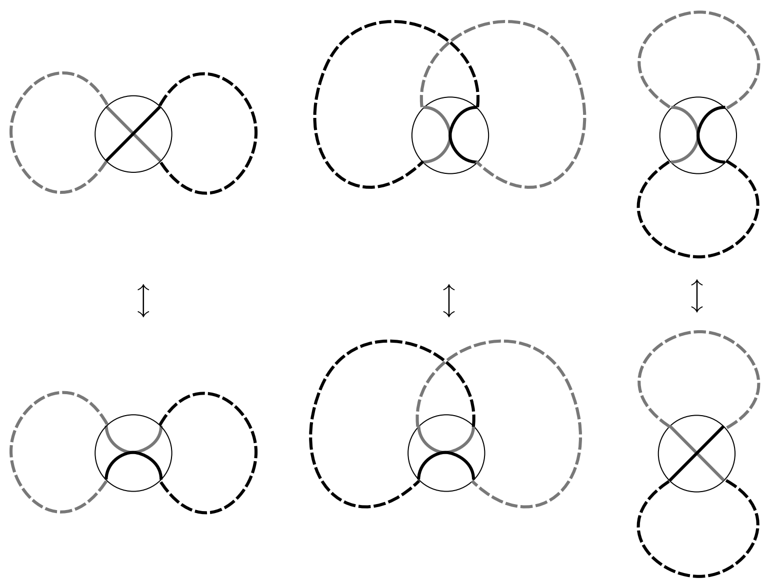

Figure 2.11. The effect of the switch operation on the local matching of a double point v, given the global matching at that double point.

We will now give a formal definition of the switch operation. Let $c = (G, d, \tau)$ be a curve and v a double point of c. Let the Euler tour be of the form

$$\tau \;=\; v h_1 \, \alpha \, h_2 v h_3 \, \beta h_4 v$$

where the h_i are the half-edges incident with v and α, β sequences of half-edges that correspond to the two global segments of v. Now, the **switch $\boldsymbol{c \circ v}$** of c at v is defined as $c \circ v := (G, d, \tau')$ where

$$\tau' \;=\; v h_2 \, \alpha^{-1} h_1 v h_3 \, \beta h_4 \, v$$

and α^{-1} is the reverse of α. This obviously defines a fixpoint-free involution. We also note that the switch changes only the Euler tour, both the multigraph G as well as its drawing d remain unaffected. As the pair (G, τ) is synonymous with the chord diagram $\mathcal{C}(c)$ this definition automatically defines an operation on chord diagrams. Precisely because G and d remain unaffected, this operation preserves the realizability of the chord diagram. It can be visualized as shown in Figure 2.12.

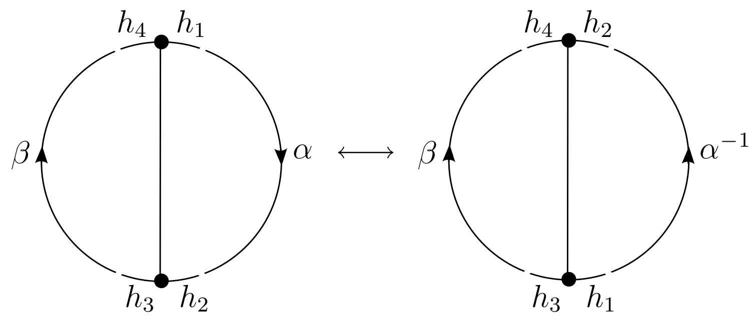

Figure 2.12. The effect of a switch on a chord diagram.

Does the switch at v change the interlacement of chords x and y?

2.3. Whether or not two given chords x and y are interlaced changes under a switch at v changes, iff both are interlaced with v.

Proof. This proof is best done by example (see Figure 2.13), but the claim also follows easily from the definition of τ'.

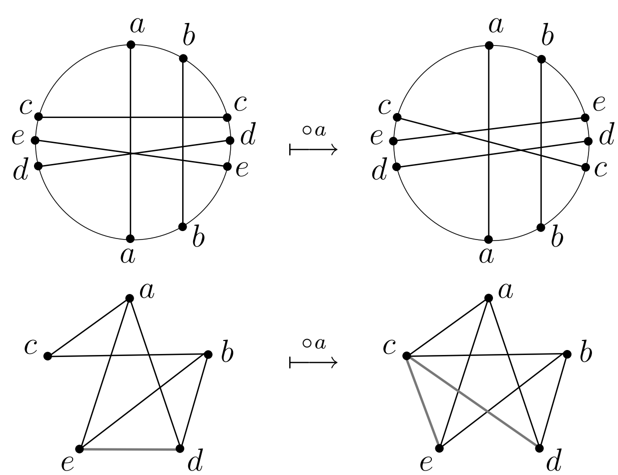

Figure 2.13. Switches and the interlacement relation: the interlacement graphs are shown below the corresponding interlacement graphs.

□

We can also formulate 2.3 in terms of interlacement graphs. Given a graph G and a a set $X \subset V(G)$, the **local complement** $G \Delta X$ of G with respect to X is the graph obtained from G by toggling every edge in $\binom{X}{2}$, i.e. it is the graph on $V(G)$ whose edge set is the symmetrical difference $E(G) \, \Delta \binom{X}{2}$.

2.4. If c is a curve with interlacement graph Λ and v a double point, then the interlacement graph $\Lambda \circ v$ of $c \circ v$ is $\Lambda \Delta N(v)$ where $N(v)$ is the neighborhood of v in c.

Motivated by 2.4 we define $\Lambda \circ v := \Lambda \Delta N(v)$. Given this last definition, we can now summarize our findings:

2.5. Switch Lemma

The switch at v, denoted $\circ v$, is an involution on curves c, chord diagrams C and interlacement graphs Λ, i.e.

$$c \circ v \circ v = c, \;\; C \circ v \circ v = C, \;\; \Lambda \circ v \circ v = \Lambda$$

Switches commute with $\mathcal{C}$ and $\mathcal{I}$, i.e.

$$\mathcal{I}(\mathcal{C}(c \circ v)) = \mathcal{I}(\mathcal{C}(c) \circ v) = \mathcal{I}(\mathcal{C}(c)) \circ v$$

Switches preserve realizability, i.e.

$$\begin{array}{rcl} C \text{ is realizable} & \Leftrightarrow & C \circ v \text{ is realizable} \\ \Lambda \text{ is realizable} & \Leftrightarrow & \Lambda \circ v \text{ is realizable} \end{array}$$

Switches do not change the underlying multigraph $\mathcal{G}(c)$, i.e.

$$\begin{array}{rcl} \mathcal{G}(c \circ v) & = & \mathcal{G}(c) \\ \mathcal{G}(C \circ v) & = & \mathcal{G}(C) \end{array}$$

There are some pitfalls though:

- switches at different chords don't commute, and
- if v is touching in c, v can be either crossing or touching in $c \circ v$, depending on the type of the diagonal matching as Figure 2.11 shows.

Note, however, that switching v has no effect on the type of another vertex w. Nonetheless, we have yet to do some work, before we fully understand the effect of the switch at v on the type of v and thus the effect of switches in general on touch-, cross- and ϑ-realizability.

2.4 Criteria for Cross-Realizability

Switches transform one curve into another while allowing us to keep track of the associated chord diagrams. In particular, given a realizable chord diagram we can obtain other realizable chord diagrams. This idea will accompany us through out the rest of this chapter and it is closely related to each of the three criteria for cross-realizability we are going to present.

Applying a switch operation to a crossing curve produces a touching double point. The Lovász-Marx criterion pursues the idea of removing touching double points that arise from the switch operation in order to preserve cross-realizability. De Fraysseix and Ossona de Mendez transform a crossing curve into a touching curve by a series of switches, because touching curve are easier to characterize than crossing curves. The Rosenstiehl criterion appears to be of an entirely different type as it does not take such an operational approach. Nonetheless, it is closely related to the switch operation as the "Rosenstiehl condition for cross-realizability" turns out to be "sort-of" an invariant of the switch operation.

Lovász-Marx Criterion

The idea is the following: From a crossing curve c, we can obtain other crossing curves with fewer double points as indicated in Figure 2.14. If a vertex v is locally crossing, both the global matching m_1 and the diagonal matching m_2 are touching. If we change the local matching at v to m_1, we obtain the drawing of two curves that touch at v as indicated on the left-hand side of Figure 2.14. Each of these two curves, when considered by itself, is crossing. We call the operation of obtaining one of these two curves from c the **loop removal**. On the other hand, if we change the local matching at v to m_2, i.e. if we switch at v, we obtain a single curve c' that is touching at v. Modifying the drawing of c' in a small neighborhood of v we can remove the touching point (as shown on the right-hand side of Figure 2.14) to, again, obtain a crossing curve. We call this operation the **switch deletion**. We thus have obtained three crossing curves that each have fewer double points than the original curve c, so that, in a sense, we can consider them to be substructures of c.

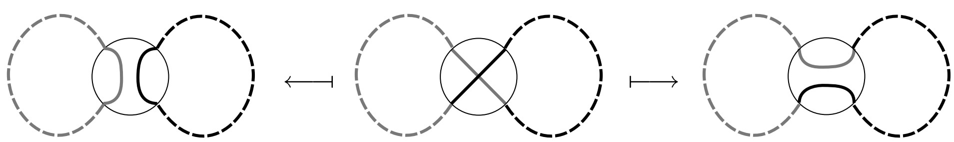

Figure 2.14. Loop removal (left) and switch deletion (right).

What operations on chord diagrams do the constructions shown in Figure 2.14 induce? Let C be the chord diagram of the curve c. The chord diagrams of the two curves obtained via loop removal are given by the global segments of C at v, respectively. The chords connecting the one segment with the other correspond to the points where the two curves meet. The chords that have both end-points in the same segment correspond to double points of the corresponding curve.

The chord diagram of the curve given by the switch deletion at v can be obtained by deleting the chord v from $C \circ v$. We use these observations to define the switch deletion and the loop removal for chord diagrams such that, by construction, both observations on chord diagrams preserve cross-realizability. Note that the effect of a loop removal on the interlacement graph cannot be described in terms of the interlacement graph alone, as the interlacement graph does not contain information about which pairs of vertices v_1, v_2 belong to the same global segment (see section 2.7).

We define a relation $\leqslant$ on chord diagrams as follows: $C_1 \leqslant C_2$ if C_1 can be obtained from C_2 by a (possibly empty) sequence of loop removals and switch deletions. For a cross-realizable C_2 all such C_1 are cross-realizable as well. Can we characterize the cross-realizable chord diagrams by giving a set of obstructions to cross-realizability under this relation?

We start out with the following elementary observation.

2.6. Two curves c_1 and c_2 that have finitely many points in common cross each other even many times.

Proof. The faces of c_1 can be 2-colored as given in 1.12. We traverse c_2 starting at a point x in a face of color A. Whenever we cross c_1, we change into a face of different color. If we touch c_1 the color stays the same. After completing our traversal we are back in the face of color A and hence have crossed c_1 even many times. □

From this we derive a lemma that is slightly more general than necessary right now, but which will prove useful again later on.

2.7. Let c be a curve and v a double point. The number of crossing neighbors of v is even if and only if v is globally touching.

In particular, if c is crossing then $\mathcal{I}(c)$ is even.

Here, the **neighborhood** of a chord v of a chord diagram or a double point v of a curve is defined as the neighborhood of v in the corresponding interlacement graph. We say a graph is **even** if every vertex has even degree.

Proof. Let s_1 and s_2 denote the global segments at v and let $\{h_1, h_2\}$ and $\{h_3, h_4\}$ be the corresponding pairs of half-edges that are matched by the global matching at v. Now $c_1 = s_1 \cup \{h_1, h_2\}$ and $c_2 = s_2 \cup \{h_3, h_4\}$ form closed curves that by 2.6 cross even many times. The common points of c_1 and c_2 are the common points of s_1 and s_2 plus v itself. Note that c_1 and c_2 cross at v, if and only if v is globally crossing in c. This means we have

$$\underbrace{\#\text{crossings of } c_1 \text{ and } c_2}_{\text{even}} = \#\text{crossings of } s_1 \text{ with } s_2 + \begin{cases} 1 & \text{if } v \text{ is globally crossing} \\ 0 & \text{if } v \text{ is globally touching} \end{cases}$$

and as the crossing neighbors of v are precisely the crossings of the global segments the first result follows. To see that $\mathcal{I}(c)$ has to be even if c is crossing, note that no vertex v of c can be globally crossing as all vertices are locally crossing. □

We now have obtained a necessary criterion for the cross-realizability of a curve. A consequence is that all K_{2n} cannot be the interlacement graphs of a crossing curve. For a given k, there is only one chord diagram C with $\mathcal{I}(C) = K_k$ and we denote it with C_{K_k}. Its Gauss code is $12\ldots k12\ldots k$, see Figure 2.15.

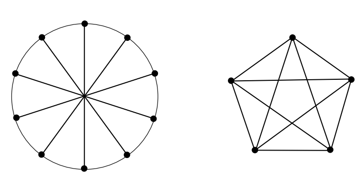

Figure 2.15. The chord diagram C_{K_5} that has K_5 as its interlacement graph.

Crossing curves realizing K_{2n+1} can be easily constructed (see Figure 2.16). We start at a point x_1 in the plane and draw a simple circle in clockwise orientation, such that x_1 becomes our first double point. We continue from x_1 in the unbounded face, following the circle in clockwise direction. We then cross into the bounded face, creating a crossing x_2. We cross back into the unbounded face at x_3 and so forth. After x_2, ..., x_{2n+1} we have crossed the circle even many times and are thus back in the unbounded face. We can now close our curve by moving to x_1 such that x_1 is locally crossing.

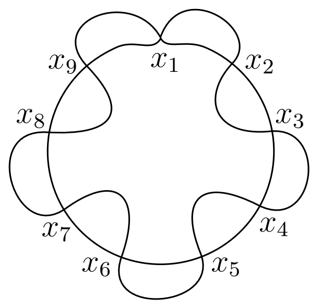

Figure 2.16. A realization of K_9 that can be generalized to a realization of any K_{2n+1}.

2.8. The interlacement graph K_i is cross-realizable if and only if i is odd.

The set $\{C_{K_{2n}}\colon n \geqslant 1\}$ presents itself as a set of obstructions to cross-realizability. The Lovász-Marx Criterion states the chord diagrams not having any $C_{K_{2n}}$ as $\leqslant$-minor are precisely the cross-realizable chord diagrams. We have already seen that this is necessary.

2.9. Lovász-Marx Criterion for Cross-Realizability

A chord diagram C is cross-realizable if and only if $C_{K_{2n}} \nleqslant C$ for all $n \geqslant 1$.

We will give a proof of 2.9 in section 2.7. Note that the set of obstructions $\{C_{K_{2n}}\colon n \geqslant 1\}$ is minimal with respect to inclusion: loop removal at any vertex of $C_{K_{2n}}$ yields the empty chord diagram which is cross-realizable and switch deletion yields the chord diagram shown in 2.17b) which can be cross-realized as shown in Figure 2.17c).

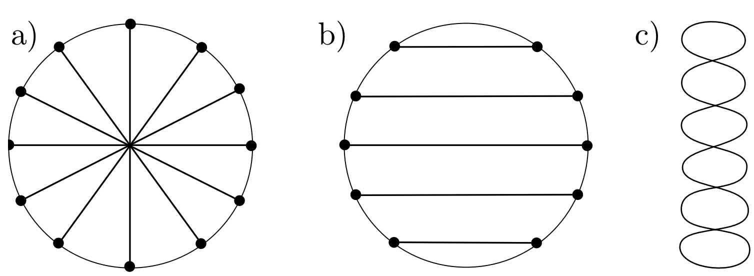

Figure 2.17. a) C_{K_6} b) $C_{K_6} \circ x \setminus x$ for any chord x. c) A realization of $C_{K_6} \circ x \setminus x$.

De Fraysseix-Ossona de Mendez Criterion

Switching a crossing vertex v turns is into a touching vertex. The types of the local matchings of all other vertices of course remain unaffected by the switch. Hence, by successively switching all crossing vertices $v_1, \dots, v_k$ we can convert a curve c into a touching curve $c \circ v_1 \circ \dots \circ v_k$ with the same number of double points. Fortunately it is easy to characterize touching curves.

2.10. Characterization of Touching Curves DE FRAYSSEIX, OSSONA DE MENDEZ

A chord diagram C is touch-realizable if and only if $\mathcal{I}(C)$ is bipartite.

Proof. In essence, the proof is given by Figure 2.18.

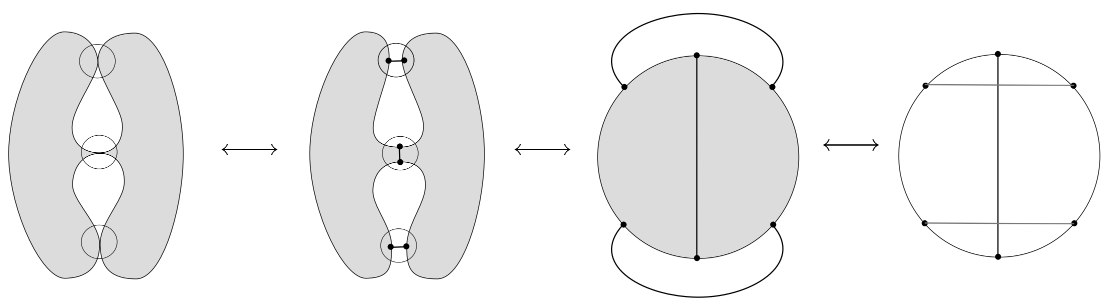

Figure 2.18. Constructing a bipartite chord diagram from a touching curve and vice versa.

Let C be a chord diagram with a bipartite interlacement graph Λ. Let (A, B) be a bipartition of the chords. Draw C in the plane such that all chords of A are drawn inside the rim and all chords of B are drawn outside the rim, such that no two chords cross. This is possible because (A, B) is a bipartition of the interlacement graph, which means that two chords in one class are not interlaced and hence can be drawn such that they do not cross. We then contract the chords, "pulling" the rim along. After all the chords have been contracted, the drawing of the rim is a touching curve with C as chord diagram.

Let c be a touching curve with chord diagram C and interlacement graph Λ. By 1.12 we can color the faces of c with black and white such that no two faces of the same color have a common curve segment in their boundary. We partition the touching points of c into classes "black" and "white" as shown in Figure 2.19.

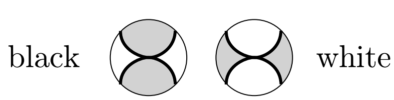

Figure 2.19. Black and white touching points.

In a small neighborhood of every touching point v, we move the two curve segments meeting at v apart and insert a chord between them. At every double point we thereby join two faces of the same color. The chords of the black touching points now lie in the white faces and the chords of the white touching points lie in the black faces. As the modified curve is now simple, it divides the plane into two faces, and since we only joined faces of the same color, the two faces we are left with are black and white. The modified curve and the chords taken together form the chord diagram of c and no two chords of the same color cross. Therefore the coloring gives us a bipartition of Λ. □

This immediately yields a criterion for the realizability chord diagrams and interlacement graphs.

2.11. Criterion for the Realizability of Chord Diagrams

A chord diagram C with $\Lambda = \mathcal{I}(C)$ is realizable, if and only if there exists some sequence of vertices $v_1, \ldots, v_k$ such that $\Lambda \circ v_1 \circ \ldots \circ v_k$ is bipartite.

Proof. Suppose c is a realization of C and $v_1, \ldots, v_k$ its crossing vertices, in any order. We now switch at these vertices to obtain a curve $c' = c \circ v_1 \circ \ldots \circ v_k$. This curve c' is touching because a switch at a crossing vertex v turns v into a touching vertex and has no effect on the local matching of any other vertex w. By 2.10 $\Lambda(c') = \Lambda \circ v_1 \circ \ldots \circ v_k$ is bipartite.

Conversely, let C, Λ and $v_1, \ldots, v_k$ be given such that $\Lambda \circ v_1 \circ \ldots \circ v_k$ is bipartite. Let c' be a touching realization of $C \circ v_1 \circ \ldots \circ v_k$. Define $c = c' \circ v_k \circ \ldots \circ v_1$. Because switches are involutions and preserve realizability we have $\mathcal{C}(c) = C$ and we are done. Note that $v_1, \ldots, v_k$ do *not* have to be crossing in c. □

By the same argument we immediately get a necessary criterion for cross-realizability, which, unfortunately, is not sufficient.

2.12. Let C be a chord diagram with chords $v_1, \ldots, v_n$.

$$C \text{ is cross-realizable} \;\Rightarrow\; C \circ v_1 \circ \ldots \circ v_n \text{ is touch-realizable} \tag{2.1}$$
$$C \text{ is cross-realizable} \;\not\Leftarrow\; C \circ v_1 \circ \ldots \circ v_n \text{ is touch-realizable} \tag{2.2}$$

Proof. We have already seen that the touch-realizability of $C \circ v_1 \circ \ldots \circ v_n$ is necessary for the cross-realizability of C. However, this is not sufficient as the example in Figure 2.20 shows. The chord diagram $C = xyxy$ is touch realizable because $\mathcal{I}(C) = K_2$ is bipartite, but $C \circ y \circ x = xyxy = C$ is not cross realizable as we know from 2.8.

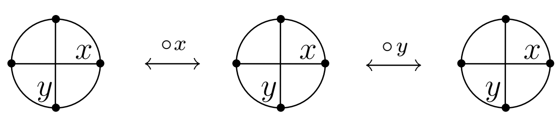

Figure 2.20. $C \circ y \circ x = C \circ y = C$

□

The proof of 2.12 is complete, but some further insight as to what is going on is desirable. Suppose we pick a touch-realization c' of $C \circ y \circ x$. We can of course compute a curve $c' \circ x \circ y$ with $\mathcal{C}(c' \circ x \circ y) = \mathcal{C}(c') \circ x \circ y = (C \circ y \circ x) \circ x \circ y = C$, so C is realizable. The problem is that we do not know which double points of $c' \circ y \circ x$ are crossing and which are touching. Consider the example in Figure 2.21. On the left we have such a touch realization c' of $C \circ y \circ x$ and switching x we obtain the curve in the middle in which x is crossing and y is touching. However, the global matching of y is crossing and its diagonal matching is touching. Hence, if we now switch y, the case in the center of Figure 2.11 applies and y remains touching in the curve on the right hand side. (Recall that a switch at v interchanges the local and the diagonal matchings of v.)

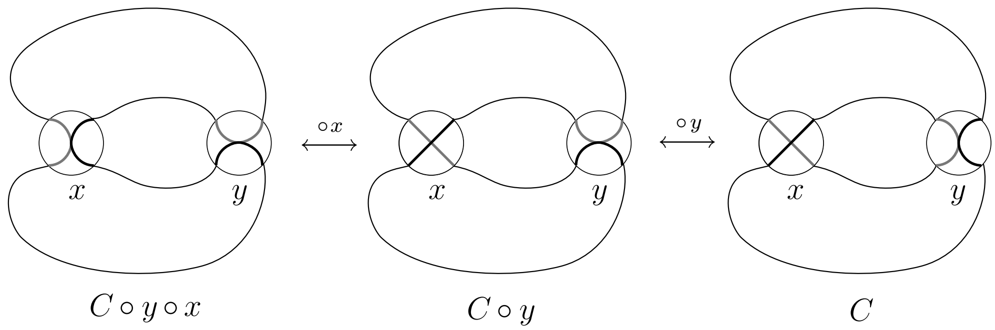

Figure 2.21. Switching the double points of a touching curve does not always yield a crossing curve.

The problem therefore is that given a chord diagram C and type function ϑ telling us which chords are supposed to be (locally) crossing and which are supposed to be (locally) touching, we do not know what the types of the diagonal matchings of the touching double points of a ϑ-realization of C are. Hence, we have no control over how the types of the touching double points change under the switch operation. To resolve this, we will introduce the notion of augmented chord diagrams in section 2.5. But first, we will complete our survey of the three characterizations of cross-realizability we want to present.

We know from 2.7 that a cross-realizable chord diagram necessarily has an even interlacement graph. The chord diagram from Figure 2.20, however, does not have an even interlacement graph. It turns out that combining the necessary conditions from 2.7 and 2.12 with 2.10 yields a sufficient condition. We will see that 2.13 is sufficient during our study of augmented chord diagrams.

2.13. De Fraysseix-Ossona de Mendez Criterion for Cross-Realizability

Let C be a chord diagram with interlacement graph Λ and vertices $v_1, \ldots, v_n$. C cross-realizable if and only if Λ is even and $\Lambda \circ v_1 \circ \ldots \circ v_n$ is bipartite.

This formulation of the de Fraysseix-Ossona de Mendez Criterion is due to Godsil and Royle [5]. In [3] de Fraysseix and Ossona de Mendez used a different form of switch operation to formulate their theorem.

Note that if we want to check the cross-realizability of a given chord diagram C, both the Lovász-Marx and the de Fraysseix-Ossona de Mendez characterizations require us to compute switches of C or respectively of its interlacement graph Λ. This is easy enough for a concrete diagram C or a concrete graph Λ. But if we do not have complete information about these objects, that is if we do not know whether two chords x, y are interlaced in C (whether they are adjacent in Λ) or not, it becomes very difficult to deduce something about the structure of $C \circ v_1 \circ \ldots \circ v_k$ or $\Lambda \circ v_1 \circ \ldots \circ v_k$. The reason is that the effect of switching v_i is determined by $N(v_i)$, but switching at another vertex v_j may change $N(v_i)$.

These observations make the Rosenstiehl Criterion of interest, because it does not rely on any form of switch operation.

Rosenstiehl Criterion

Let v, w denote vertices of an interlacement graph Λ. If Λ is the interlacement graph of a crossing curve, all vertices in the common neighborhood of $N(v) \cap N(w)$ are crossing. The parity of the common crossing neighborhood is a quantity that will be of great interest to us. We a say a pair of vertices v, w is **even** if their common crossing neighborhood is even and we say v, w are **odd** if it is odd. In particular we can apply these terms to edges of Λ.

2.14. Rosenstiehl Criterion for Cross-Realizability

An interlacement graph Λ is cross-realizable, if and only if all of the following conditions hold:

i. Λ is even.

ii. If vertices v, w of Λ are not interlaced, the pair v, w is even.

iii. The set of even edges forms a cut in Λ.

There are several things to note about this criterion. First of all, if we let $v = w$ then ii. reduces to the statement that Λ is even, so condition i. is redundant. We listed i. separately, because it is easiest to digest and we have already seen that it is necessary. Using a similar argument we can show that ii. is necessary as well.

Proof that 2.14.ii. is necessary. If $v \neq w$ are not interlaced in c, one global segment of v does not meet w and vice versa. In Figure 2.22 these two segments s_1, s_2 are on the left and right hand side. Using appropriate local segments at v and w, we can form two closed curves c_1 and c_2, that have to cross even many times. However, the neighbors of v are precisely the crossings of the one global segment at v with the other, so the common neighbors of v and w are simply the crossings of s_1 with s_2.

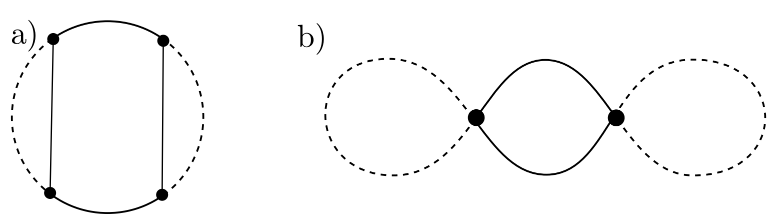

Figure 2.22. The common neighbors of two non-interlaced chords are precisely the points the two dashed segments of the curve have in common.

For vertices $v = w$ lemma 2.7 states that ii. is necessary. □

As we have seen the essence of i. and ii. is the fact 2.6 that two curves cross each other even many times. However, this argument cannot be applied if v and w are interlaced. To understand iii. we will have to make use of a different geometric insight.

iii. states that the vertices of Λ can be partitioned into two classes A, B such that two interlaced vertices v, w are of the same class if and only if their common crossing neighborhood is odd. This is noteworthy insofar as our criterion for touch-realizability also required the existence of a bipartition of the vertex set. Instead of giving a proof that iii. is necessary, we will provide a geometric interpretation of this bipartition in the case of crossing curves.

Again the bipartition of the double points of a curve is directly related to the bipartition of its faces given by a proper coloring of the faces with black and white. Consider Figure 2.23. a) shows an example of a crossing curve. b) shows the corresponding chord diagram along with the two classes of chords as given by the Rosenstiehl Criterion. c) shows the interlacement graph in which the even edges form a cut. Let us now single out the double points marked x and y. Small neighborhoods around each are shown in d) and e), respectively, along with the corresponding global matching. The half-edges matched by the global matching at x "share" a black face while the half-edges matched by the global matching at y "share" a white face. This is not an accident: it will turn out that this is precisely what characterizes the two classes in the Rosenstiehl bipartition. In the next section we will introduce the concept of an augmented chord diagram to capture this difference between x and y.

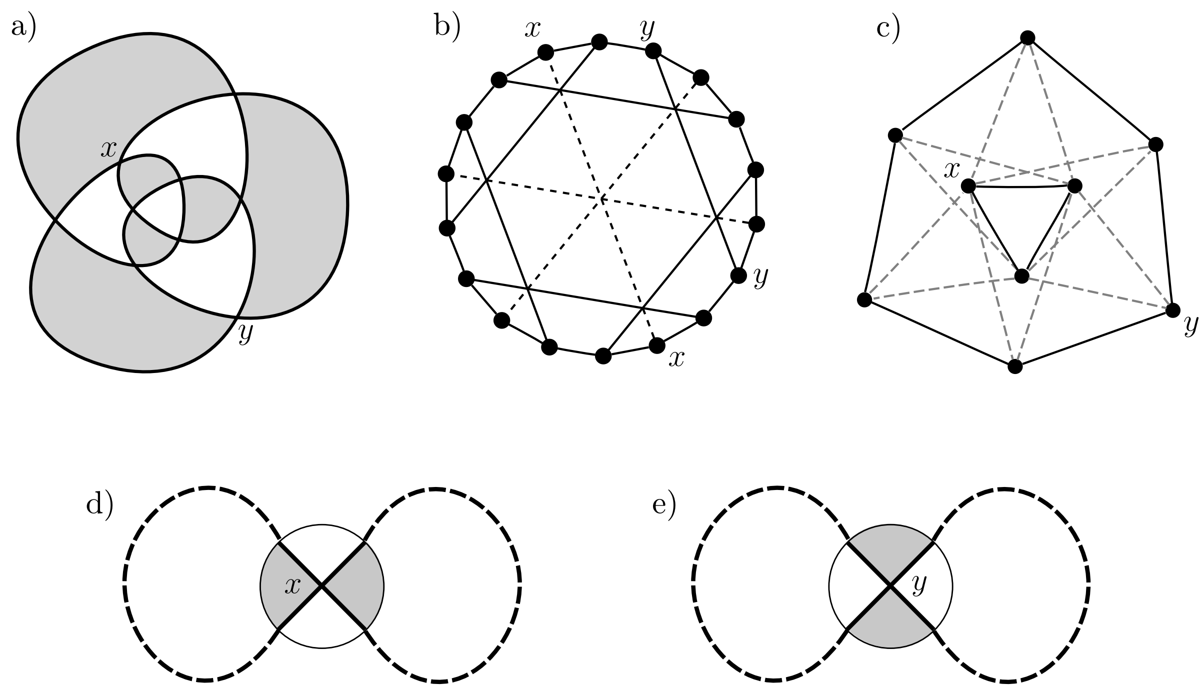

Figure 2.23. a) A crossing curve c. b) The chord diagram $\mathcal{C}(c)$. The chords belonging to the one class of the Rosenstiehl bipartition are dashed, and those belonging to the other class are solid. c) The interlacement graph $\mathcal{I}(c)$. Even edges are dashed. d) The local and global matching at x along with the colors of the faces in a small neighborhood of x. The faces between the half-edges matched by the global matching are black. e) The local and global matching at y. The faces between the half-edges matched by the global matching are white.

A few words on the history of the Rosenstiehl Criterion. Rosenstiehl's original proof [9] of 2.14 builds on the theory of maps. He also gives the geometric interpretation of the bipartition mentioned above, although he presents it in an entirely different fashion. In [3] de Fraysseix-Ossona de Mendez give a simpler proof of 2.14 that is based on (a variant of) the switch operation. Following their approach and using our new concept of augmented chord diagrams, we will now proceed as follows: In section 2.5 we will introduce augmented chord diagrams and give a criterion, 2.20, for their realizability in terms of the switch operation. This is a generalization of the de Fraysseix-Ossona de Mendez Criterion insofar as it characterizes curves that can have both touching and/or crossing double points. In section 2.6 we will then prove corresponding generalization of Rosenstiehl's theorem. In section 2.7 we derive the Lovász-Marx Criterion from the Rosenstiehl Criterion as given in [1]. In preparation for section 2.6 we will now give a very compact formulation of the Rosenstiehl Criterion that is also from [3]. By our above observations it is immediate that 2.14 and 2.15 are equivalent.

2.15. Rosenstiehl Criterion – Compact Version

An interlacement graph Λ is cross-realizable, if and only if there exists a bipartition (A, B) of the vertices such that for all vertices v, w the following holds:

$$\text{the pair } v, w \text{ is odd} \iff v, w \text{ are interlaced and of the same class}$$

2.5 Augmented Chord Diagrams

Let us summarize some of our findings so far.

- A curve can be interpreted as a multigraph G, with a rotation system π and an Euler tour τ. Chord diagrams correspond to pairs (G, τ). We have yet to find a fitting representation of π in the context of chord diagrams.

- A chord diagram C and correspondingly a pair (G, τ) define a local, a global, and a diagonal matching at every double point v.

- The switch operation turns crossing vertices into touching vertices, but we have no control over when a touching vertex becomes crossing. This is the last obstacle to a proof of the de Fraysseix-Ossona de Mendez Criterion for cross-realizability

- The faces of a curve can be 2-colored. This coloring partitions the touching vertices as well as the crossing vertices of a curve into two classes. These bipartitions play a central role in the de Fraysseix-Ossona de Mendez and Rosenstiehl Criteria for touch- and cross-realizability.

We are now going to put all of the above together by defining a new object called the augmented chord diagram.

The faces of c can be colored with black and white. Given such a 2-coloring, the rule shown in Figure 2.24 defines a 3-coloring of the matchings at every vertex with the colors “black”, “white” and “crossing”. Changing the 2-coloring of the plane swaps the colors black and white, therefore we identify colorings in which black and white have been interchanged at every vertex. The 3-coloring of the matchings is then uniquely determined by the curve c. More precisely it is uniquely determined by the induced graph G and its embedding d in the plane. How c traverses G is irrelevant.

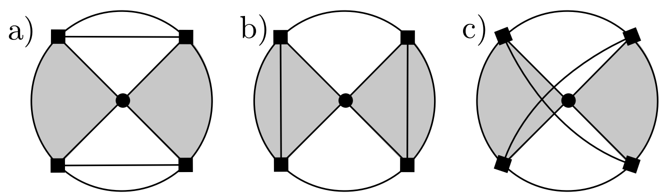

Figure 2.24. The a) white, b) black and c) crossing matching.

A curve c thus defines two 3-colorings of the 3 perfect matchings at every double point. The one with the colors “local”, “global” and “diagonal” is given the interlacement properties of c, i.e. by the multigraph G and the Euler tour τ. The one with colors “black”, “white” and “crossing” is given by the geometry of the curve, i.e. by the multigraph G and its rotation system π. The pair $(G,\,\tau)$ is determined by the chord diagram C while π is not. We now augment C with information about the geometry of c.

We denote the set of perfect matchings on th half-edges incident with a double point v by M_v. An **augmented chord diagram** $(C,\,(a_v)_v)$ is a chord diagram C together with a family $(a_v)_{v\text{ a chord of }C}$ of bijections

$$a_v\colon M_v \to \{\text{“crossing”, “black”, “white”}\}$$

Here we identify families $(a_v)_v$ and $(s\circ a_v)_v$, where s is the bijection $s\colon\{$“crossing”, “black”, “white”$\}\to\{$“crossing”, “black”, “white”$\}$ that swaps black and white – just as we identified rotation systems $(\pi_v)_v$ and $(\pi_v^{-1})_v$. Note that as C defines a bijection between M_v and $\{$“local”, “global”, “diagonal”$\}$, we can equivalently regard $(a_v)_v$ as a family of bijections

$$a_v\colon \{\text{“local”, “global”, “diagonal”}\} \;\to\; \{\text{“crossing”, “black”, “white”}\}$$

and we will use both representations interchangeably. We call the family $a=(a_v)_v$ as well as the pair $(C,\,a)$ an **augmentation** of C and we write $\mathcal{A}(c)$ for the augmented chord diagram of a curve c.

2.16. The augmented chord diagram $\mathcal{A}(c)$ of a curve c is well defined.

Proof. The (equivalence class of a) curve c has a unique rotation system, which in turn defines the dual graph $\mathcal{G}(c)^*$ uniquely. $\mathcal{G}(c)^*$ is bipartite and connected, and hence the bipartition (A, B) of $\mathcal{G}(c)^*$ is uniquely determined up to the identification of (A, B) with (B, A) which corresponds to our identification of $(a_v)_v$ with $(s \circ a_v)_v$. ☐

The augmented chord diagram of the example in Figure 2.2 is given in Figure 2.25.

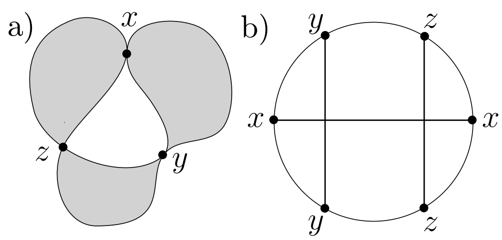

c)

	x	y	z
local	black	white	crossing
global	crossing	black	black
diagonal	white	crossing	white

Figure 2.25. a) A curve. b) Its chord diagram. c) Its augmentation.

We have already defined what it means for a double point to be "locally crossing", "globally touching", etc. In addition we now say that a double point v is **locally black** iff $a_v(\text{local}) = \text{black}$, that v is **globally white** iff $a_v(\text{global}) = \text{white}$, etc. Thus, a double point v is crossing iff it is locally crossing, and v is touching iff it is locally black or locally white. A pair of a chord diagram C and a type function ϑ is thus augmented by $(a_v)_v$ iff $a_v(\text{local}) = \text{crossing} \Leftrightarrow \vartheta(v) = \text{crossing}$.

We say an augmented chord diagram (C, a) is **realizable**, iff there exists a curve c such that $(C, a) = \mathcal{A}(c)$. This is a much stronger concept then the concept of realizability for chord diagrams. In particular the color of the local matching is given for every double point v, which determines whether v is crossing or touching. So, given a chord diagram C, a type function ϑ and an augmentation a of the two, realizability of (C, a) implies ϑ-realizability of C but not vice versa.

Now, the augmentation a in the pair (C, a) has the role of the rotation system π in the triple (G, π, τ). We will study the relationship between augmentations and rotation systems in detail in section 2.8. In particular we will see that for every realizable augmented chord diagram (C, a) there is exactly one equivalence class c of curves realizing (C, a).

Augmented Chord Diagrams and Switches

Let us consider the effect of a switch on an augmented chord diagram. During a switch the underlying graph and its embedding remain the same, only the tour through the graph changes. Therefore the coloring given by the embedding of the graph remains the same and so we have to concentrate on the coloring given by the tour. Let A be an augmented chord diagram and x the chord we switch. As Figure 2.26 shows, the local and diagonal matchings of x are interchanged. The global matching then has to remain the same.

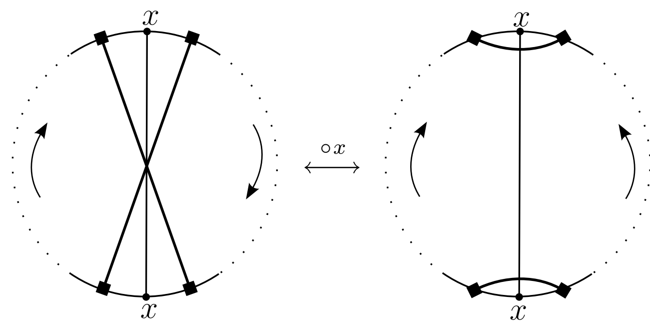

Figure 2.26. The switch at a chord x swaps the local and diagonal matchings at x.

Let y be some chord interlaced with x. As Figure 2.27 shows, the global and the diagonal matchings are interchanged. The local matching remains unchanged.

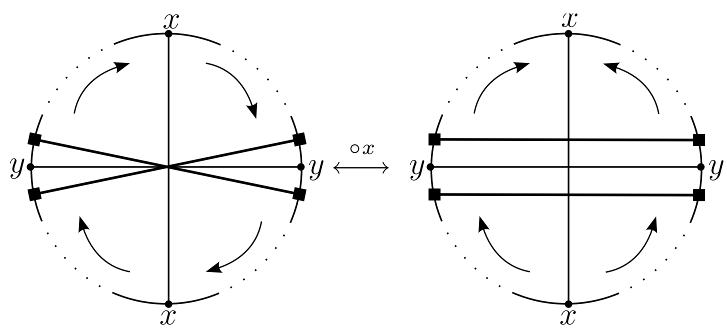

Figure 2.27. The switch at a chord x swaps the global and diagonal matchings of every chord y that is interlaced with x.

Finally, for a chord z different from but not interlaced with x, all three matchings do not change. These observations allow us to give the augmentation of $\mathcal{A}(c \circ x)$ in terms of the augmentation of $\mathcal{A}(c)$ for any curve c. If $(a_v)_v$ is the augmentation of $\mathcal{A}(c)$ the augmentation $(a'_v)_v$ of $\mathcal{A}(c \circ x)$ is given by

$$\begin{aligned} a'_v(\text{local}) &= \begin{cases} a_v(\text{diagonal}) & \text{if } x = v \\ a_v(\text{local}) & \text{otherwise} \end{cases} \\ a'_v(\text{global}) &= \begin{cases} a_v(\text{diagonal}) & \text{if } x, v \text{ are interlaced} \\ a_v(\text{global}) & \text{otherwise} \end{cases} \\ a'_v(\text{diagonal}) &= \begin{cases} a_v(\text{local}) & \text{if } x = v \\ a_v(\text{global}) & \text{if } x, v \text{ are interlaced} \\ a_v(\text{diagonal}) & \text{otherwise} \end{cases} \end{aligned}$$

For an augmented chord diagram A, we define the augmentation of $A \circ x$ to be $(a'_v)_v$. Note that this again yields an involution as can be seen from Figures 2.26 and 2.27. We can now formulate a version of the Switch Lemma for augmented chord diagrams, which, although trivial to prove, is more powerful than the original Switch Lemma, because the concept of realizability is much more restrictive for augmented chord diagrams.

2.17. Augmented Switch Lemma

The switch at a double point v is an involution on the augmented chord diagrams, switches commute with $\mathcal{A}$ and

$$A \text{ is realizable } \Leftrightarrow A \circ v \text{ is realizable}$$

Proof. Because $\circ$ is an involution on augmented chord diagrams, we only have to show that if A is realizable, so is $A \circ v$. Let c be a realization of A. By construction $\mathcal{A}(c \circ v) = \mathcal{A}(c) \circ v = A \circ v$, so $c \circ v$ is a realization of $A \circ v$. □

Let C be a chord diagram with chords $v_1, \ldots, v_k$ and $A = (C, (a_v)_v)$ an augmentation. Suppose A is locally crossing at all vertices. Then $A \circ v_k \circ \ldots \circ v_1$ is locally touching at all vertices and it is a consequence of 2.17 that

$$A \circ v_k \circ \ldots \circ v_1 \text{ is realizable } \Leftrightarrow A \text{ is realizable} \tag{2.3}$$

Recall from 2.12 that this is not true for "ordinary" chord diagrams. We proved 2.12 by applying our criteria for touch- and cross-realizability to the chord diagram $C = xyxy$. To illustrate what is going on with an example (see Figure 2.21) we picked a touch realization c' of $C \circ y \circ x$ and noted that when passing from $c' \circ x$ to $c' \circ x \circ y$, the chord y does not become locally crossing because in $c' \circ x$ it is globally crossing. Switching y only swaps the local and diagonal matchings which are both touching.

What can we learn about this example from the point of view of augmented chord diagrams? In general, we have for any augmentation A of $C = xyxy$, that

$$\begin{aligned} a_y^{A\circ y\circ x}(\text{global}) &= a_y^{A\circ y}(\text{diagonal}) = a_y^{A}(\text{local}) \\ a_y^{A\circ y\circ x}(\text{diagonal}) &= a_y^{A\circ y}(\text{global}) = a_y^{A}(\text{global}) \\ a_y^{A\circ y\circ x}(\text{local}) &= a_y^{A\circ y}(\text{local}) = a_y^{A}(\text{diagonal}) \end{aligned}$$

We now claim:

2.18. A touching curve c is globally touching at every double point v.

Proof. We modify the local matching so that it equals the global matching, thus obtaining two curves c_1 and c_2. These two do not cross outside v as c is touching. Hence, because two curves cross each other even many times (cf. 2.10), they do not cross at v. □

To see that the example $C = xyxy$ is not cross-realizable can proceed in several ways: Keeping track of the augmentation of y we can argue that given a cross-realization of C, y has to be globally crossing in $C \circ y \circ x$, so by 2.18 $C \circ y \circ x$ cannot be touch-realizable. We could also argue that by 2.18 y has to be diagonally crossing in a touch-realization of $C \circ y \circ x$ and hence y will be globally crossing in C and not locally crossing as desired.

Realizability

2.10 and 2.18 give us necessary criteria for the realizability of an augmented chord diagram that is locally touching everywhere. Taking these two conditions together we obtain a characterization.

2.19. Realizability of Touching Augmented Chord Diagrams

An augmented chord diagram that is locally touching at every chord is realizable if and only if

$$\left.\begin{array}{l}\text{no two chords with the same local color are interlaced, and}\\ \text{every chord is diagonally crossing.}\end{array}\right\} \qquad (2.4)$$

Proof. Suppose A is a realizable augmented chord diagram that is locally touching everywhere. Let c be a realization of A. From 2.10 we know that $\mathcal{I}(c)$ is bipartite, but from the proof of 2.10 we know that the bipartition is given by the color of the local matchings. So no two chords with the same local color can cross. From 2.18 we know that every chord is diagonally crossing.

Given an augmented chord diagram A that the meets the conditions, we start by drawing the chord diagram in the plane, coloring the inside of the circle black and flipping all the locally black chords to the outside. Contracting all the chords gives us a touching curve in which all the local matchings have the correct color. We still have to see that the other matchings have the color given by A. Because, by construction, we have created a touching curve, all its diagonal matchings are crossing as required. Given the local and diagonal matchings at every chord, the global matching is uniquely determined. □

They key point here is that if A had a chord that was not diagonally crossing, the construction from the proofs of 2.10 and 2.19 simply would not yield a curve with the correct diagonal matching.

Note that the bipartition mentioned in 2.10 is given by the colors of the local matchings. So if we are given an augmented chord diagram, 2.19 is easier to check than 2.10, because we do not have to look for a bipartition, we only have to check the given one. If we are only given the chord diagram and the constraint that all chords are supposed to be touching, finding an augmentation as required in 2.19 may seem more difficult than finding a bipartition as in 2.10. However, this is not the case: finding a bipartition is equivalent to finding an augmentation, because the diagonal matching is given by (2.4). A bipartition defines the color of the local matching and the global matching is then uniquely determined.

We can now provide an algorithmic criterion for the realizability of arbitrary augmented chord diagrams. Using augmented chord diagrams we can prescribe for every chord individually whether it should be touching or not.

2.20. Realizability of Augmented Chord Diagrams – Algorithmic Version

Let A be an augmented chord diagram and v_1, ..., v_k the locally crossing chords of A.

$$A \text{ is realizable } \Leftrightarrow \text{ (2.4) holds for } A \circ v_1 \circ \ldots \circ v_k$$

Proof. By 2.17, A is realizable iff $A \circ v_1 \circ \ldots \circ v_k$ is realizable. By 2.19, $A \circ v_1 \circ \ldots \circ v_k$ is realizable iff (2.4) holds. □

Now, what does this criterion for the realizability of augmented chord diagrams tell us about the ϑ-realizability of chord diagrams? To find out whether a given chord diagram C and is ϑ-realizable for a given type function ϑ, we could, of course, enumerate all augmentations A of (C, ϑ) and use 2.20 to check if any of those is realizable. But as there are more than $2^{\#\text{chords}}$ augmentations of (C, ϑ) this would be horribly inefficient. The algorithm to apply is the following: We are given C and ϑ. If (C, ϑ) is realizable we want to compute an augmentation $(a_v)_v$ realizing (C, ϑ). If (C, ϑ) is not realizable, we want to obtain a proof of this fact. To that end, we introduce variables a_v^m for each chord v and each $m \in \{\text{local}, \text{global}, \text{diagonal}\}$ and proceed as follows:

1. Let v_1, ..., v_k denote the chords with $\vartheta(v_i) = \text{crossing}$, in any order. Assign $a_{v_i}^{\text{local}} := \text{crossing}$ for every such v_i.

2. Put $a_v(m) = a_v^m$ for every v and m.

3. Compute $(C', (a'_v)_v) = (C, (a_v)_v) \circ v_1 \ldots \circ v_k$.

4. If there is a chord v with $a'_v(\text{global}) = \text{crossing}$, C is not ϑ-realizable. Otherwise:

   - For every v with $\vartheta(v) = \text{crossing}$, we have $a'_v(\text{diagonal}) = \text{crossing}$.

   - For all other vertices w, we have $a'_w(\text{diagonal}) = a_w^m$ for some m and assign $a_w^m := \text{crossing}$.

5. Compute a bipartition of the chords of C', i.e. a 2-coloring in which chords of the same color are not interlaced.

   - If such a bipartition does not exist, C is not ϑ-realizable.

   - Otherwise, let A, B be such a bipartition. For every vertex v we have $a'_v(\text{local}) = a_v^m$ for some m. We now assign

$$a_v^m := \begin{cases} \text{black}, & \text{if } v \in A \\ \text{white}, & \text{if } v \in B \end{cases}$$

6. For each v, we have at this point assigned values to exactly two of the variables a_v^{local}, a_v^{global}, a_v^{diagonal} and these two values differ. We now assign the third value to the third variable. This defines an augmentation of C that is ϑ-realizable.

Note that any realizable augmentation of (C, ϑ) uniquely determines the rotation system and thus the equivalence class of a realization. See section 2.8 for details.

The above algorithm also has one important theoretical consequence:

2.21. Let C be a chord diagram, ϑ a type function and v_1, ..., v_k the crossing chords of (C, ϑ). $C \circ v_1 \ldots \circ v_k$ is bipartite if and only if there exists an augmentation A of (C, ϑ) with the property that in $A \circ v_1 \ldots \circ v_k$ no two chords of the same color are interlaced.

Proof. If there exists such an augmentation, then a bipartition of $C \circ v_1 \ldots \circ v_k$ is given by the colors of the local matchings in $A \circ v_1 \ldots \circ v_k$. Conversely, given a bipartition of $C \circ v_1 \ldots \circ v_k$ we simply run the algorithm, do not abort in step 4. even if we encounter a globally crossing vertex, and use the bipartition defined in step 5. □

Now, if we interpret 2.20 as a criterion for cross-realizability, it is very similar to the de Fraysseix-Ossona de Mendez Criterion for cross-realizability as formulated by Godsil and Royle (i.e. 2.13). The difference is that in 2.20 the augmented chord diagram $A \circ v_1 \ldots \circ v_n$ is required to be diagonally crossing everywhere, while in 2.13 Λ is required to be even. Therefore, can we give derive 2.13 from 2.20?

How can we interpret the requirement that every chord of $A \circ v_1 \ldots \circ v_k$ is diagonally crossing? We observe that v_i is locally crossing in each of the augmented chord diagrams $A, \ldots, A \circ v_1 \ldots \circ v_{i-1}$ and it becomes diagonally crossing in $A \circ v_1 \ldots \circ v_i$. After that, we have for any $j > i$ that the global and diagonal matchings at v_i are interchanged in the step from $A \circ v_1 \ldots \circ v_{j-1}$ to $A \circ v_1 \ldots \circ v_j$ if and only if v_i and v_j are interlaced in $A \circ v_1 \ldots \circ v_{j-1}$. We call a step with this property an inversion of v_i. Now, v_i is diagonally crossing in $A \circ v_1 \ldots \circ v_k$ if and only if v_i is inverted even many times. All we need to complete the proof of 2.13 is to show that all chords v_i are inverted even many times if Λ is even. We will postpone this proof until the end of section 2.6, though, since the material developed there will turn out to be useful for this purpose.

We will now seek to develop a Rosenstiehl-type criterion for augmented chord diagrams.

2.6 A Combinatorial Characterization

We have yet to define a notion of "class" for curves with locally crossing double points. If x is touching, the **class** is given by the color of the local matching. What if x is crossing? As motivated in section 2.4, we define the **class** of a locally crossing double point x to be the color of the diagonal matching. Recall that the colors of the local and diagonal matchings of x are interchanged by a switch at x, so the class of x does not change under a switch at x, unless x is globally crossing. Given this definition, the class of every double point is either "black" or "white".

We can now formulate a Rosenstiehl-type criterion that applies to arbitrary curves, i.e. curves with crossing and/or touching double points. Again, we call a chord x **even** iff the number of crossing neighbors is even and **odd** otherwise. We call a pair x, y of chords **even** iff the number of common crossing neighbors is even, and **odd** otherwise.

2.22. Realizability of Augmented Chord Diagrams – Combinatorial Version

An augmented chord diagram A is realizable if and only if both of the following conditions hold:

i. A chord x is odd, if and only if x is globally crossing.

ii. Two chords $x \neq y$ are odd, if and only if x, y are interlaced and of the same class.

Note that both the Rosenstiehl Criterion 2.14 (resp. 2.15) for cross-realizability of chord diagrams as well as our criterion 2.19 for the realizability of touching augmented chord diagrams are special cases of 2.22.

Proof of 2.19 from 2.22. If A is touching, the number of crossing neighbors of x is always zero and hence even, so i. reduces to the condition that all chords are diagonally crossing. By the same argument ii. reduces to the condition that no chords x and y are interlaced and of the same class. These are just the conditions of 2.19. □

We will not be able to obtain 2.19 in this fashion, though, as we are going to make use of 2.19 in our proof of 2.22. However, the above motivates 2.22 nicely. The derivation of the Rosenstiehl Criterion from 2.22, on the other hand, is of greater consequence, as we have not given a proof of that one yet.

Proof of 2.14 from 2.22. If A is crossing, i. simply states that every chord x has even many neighbors which is part of the Rosenstiehl Criterion. Condition ii. reduces to the statement that two chords x and y have an odd neighborhood, if and only if they are interlaced and of the same class. If we split this statement by regarding interlaced and non-interlaced pairs of chords separately, we obtain precisely the Rosenstiehl conditions: For chords x and y that are not interlaced, the common neighborhood is even. Interlaced chords on the other hand have an odd common neighborhood, if and only if they are of the same class, which means that the even edges form a cut in the interlacement graph.

One fine point here is that the Rosenstiehl Criterion requires the existence of *some* partition of the chords into two classes with the given properties, while 2.22 requires that the specific partition defined at the beginning of this section has the given properties. The difference is of course that 2.22 talks about augmented chord diagrams, while the Rosenstiehl Criterion does not. If we want to show that the Rosenstiehl condition is sufficient using 2.22, we start out with a chord diagram C that meets the condition. The even edges form a cut in $\mathcal{I}(C)$ which defines a bipartition X, Y of the set of chords. An augmentation $(a_v)_v$ of C is then defined by $a_v(\text{local}) = \text{crossing}$ and

	a_v(diagonal)	a_v(global)
$v \in X$	black	white
$v \in Y$	white	black

for each chord v. The augmented chord diagram $(C,\,(a_v)_v)$ then meets the condition of 2.22 as described.

If we want to show that the Rosenstiehl condition is necessary, we obtain an augmented chord diagram $(C,\,(a_v)_v)$ from the given curve, and by 2.22 we know that the even edges form a cut, whose bipartition of the set of chords is given by the diagonally black and the diagonally white double points. □

Our criterion for the realizability of augmented chord diagrams thus provides us explicitly with a combinatorial interpretation of the bipartition of the set of chords mentioned in the Rosenstiehl characterization of crossing curves. Going through Rosenstiehl's original proof of his theorem carefully yields the same interpretation.

We will now turn to the proof of 2.22. It consists of these four steps:

1. The theorem holds for touching augmented chord diagrams as we have already seen.
2. Every augmented chord diagram A can be converted into a touching augmented chord diagram by a sequence of switches $\circ\, v_1 \ldots \circ\, v_k$.
3. We will show that the condition of 2.22 is invariant under switches.
4. We complete the proof as follows:

$$\begin{array}{rl} & A \text{ is realizable} \\ \Leftrightarrow & A \circ v_1 \ldots \circ v_k \text{ is realizable} \\ \Leftrightarrow & \text{the condition of 2.22 holds for } A \circ v_1 \ldots \circ v_k \\ \Leftrightarrow & \text{the condition of 2.22 holds for } A \end{array}$$

This is the approach taken by de Fraysseix and Ossona de Mendez in [3] in their proof of Rosenstiehl's theorem. As in their proof, showing the invariance of our condition under switches will be simply a matter of calculation. In this more general setting, however, the calculation will be more involved.

We will first write the condition of 2.22 as a system of equations $P(C) \equiv 0$ over $\mathbb{Z}_2$, which is in essence a boolean expression. To that end we will associate with a given augmented chord diagram A a set of boolean predicates. Each is 1 if the statement in the square brackets is true and 0 if it is false. For chords x, y we define

$$\begin{array}{rcl} g_x &=& [a_x(\text{global}) = \text{crossing}] \\ \ell_x &=& [a_x(\text{local}) = \text{crossing}] \\ c_x &=& [\text{the class of } x \text{ is white}] \\ i_{x,y} &=& [x \text{ and } y \text{ are interlaced}] \\ \delta_{x,y} &=& [x = y] \end{array}$$

We will also consider the incidence vector N_x of the crossing neighborhood of x, i.e.

$$N_x = (i_{x,y} \cdot \ell_y)_{y \text{ a chord of } A}$$

Given this definition, we have

$$\langle N_x, N_y \rangle := \sum_{z \text{ a chord of } A} i_{x,z} \cdot \ell_z \cdot i_{y,z} \cdot \ell_z = \#\text{common crossing neighbors of } x \text{ and } y$$

We can now write the characterization of realizable augmented chord diagrams as a system of equations.

2.23. Realizability of Augmented Chord Diagrams - Equational Version

An augmented chord diagram A is realizable if and only if for all chords x, y the following equation holds:

$$\langle N_x, N_y \rangle + i_{x,y} \cdot (c_x + c_y + 1) + \delta_{x,y} \cdot g_x \equiv 0 \tag{2.5}$$

We write $P(A) \equiv 0$ to denote that all of these equations hold for A.

We first have to see that 2.23 is equivalent to 2.22. Note that addition and multiplication over $\mathbb{Z}_2$ are simply the boolean operations "exclusive or" and "and", respectively. Hence an equation of the form $a + b \equiv 0$ is equivalent to the boolean expression $a \leftrightarrow b$ (i.e. "a is equivalent to b") and an equation of the form $a + b \equiv 1$ is equivalent to $a \dot{\vee} b$ (i.e. "either a or b"). The expression $(c_x + c_y + 1)$ in (2.5) is therefore 1 iff c_x and c_y have the same value.

Considering the cases $x = y$ and $x \neq y$ separately, we can thus reformulate (2.5) as follows. If $x = y$, x is not interlaced with y, so (2.5) reduces to

$$\langle N_x, N_x \rangle + g_x \equiv 0$$

which is precisely 2.22.i. If $x \neq y$, (2.5) reduces to

$$\langle N_x, N_y \rangle + i_{x,y} \cdot [x \text{ and } y \text{ are of the same class}] \equiv 0$$

which is precisely 2.22.ii. Our task is now to show the following lemma:

2.24. Let A denote an augmented chord diagram and x any chord of A. Then

$$P(A) \equiv 0 \;\Leftrightarrow\; P(A \circ x) \equiv 0$$

As outlined above, we can then complete our proof of 2.22 using 2.24.

Proof of 2.22. Let c be any curve and A its chord diagram. Let $v_1, \ldots, v_k$ be the crossing vertices of c. The curve $c \circ v_1 \circ \ldots \circ v_k$ and hence $A \circ v_1 \circ \ldots \circ v_k$ are realizable and by 2.19 the equation $P(A \circ v_1 \circ \ldots \circ v_k) \equiv 0$ holds. By 2.24 the equation $P(A) \equiv 0$ follows.

Conversely, let A be an augmented chord diagram for which the condition of 2.22 and thus $P(A) \equiv 0$ hold. Let $v_1, \ldots, v_k$ be the crossing vertices of A. By 2.24 we have $P(A \circ v_1 \circ \ldots \circ v_k) \equiv 0$ and by 2.19 there exists a realization c' of $A \circ v_1 \circ \ldots \circ v_k$. By the Augmented Switch Lemma $c' \circ v_k \circ \ldots \circ v_1$ is then a realization of A. □

To be able to calculate whether 2.24 holds, we need to calculate the parameters $g_x^{A \circ x}$, $i_{x,y}^{A \circ x}$, $c_x^{A \circ x}$ and $\langle N_x^{A \circ x}, N_y^{A \circ x} \rangle$ of $A \circ x$ given the corresponding parameters of A. This is the combinatorial part of the proof of 2.24.

2.25. Let A be an augmented chord diagram and let N, i, g, ℓ, δ, c be the parameters of A as defined above. We denote the parameters of $A \circ x$ by the same symbols with an upper index. Then, the following equations hold:

$$\begin{aligned}
N_y^{A\circ x} &\equiv N_y + i_{x,y} \cdot (N_x + (g_x + 1) \cdot x + \ell_y \cdot y) \\
g_y^{A\circ x} &\equiv g_y + i_{x,y} \cdot (\ell_y + 1) \\
\ell_y^{A\circ x} &\equiv \ell_y + \delta_{x,y} \cdot (g_y + 1) \\
c_y^{A\circ x} &\equiv c_y + i_{x,y} \cdot \ell_y + \delta_{x,y} \cdot g_x \\
i_{y,z}^{A\circ x} &\equiv i_{y,z} + i_{x,y} \cdot i_{x,z} \cdot (\delta_{y,z} + 1)
\end{aligned}$$

Proof of 2.25. $N_y^{A\circ x} \equiv N_y + i_{x,y} \cdot (N_x + (g_x + 1) \cdot x + \ell_y \cdot y)$

Note that in this first equation, N_x, N_y, x, y are vectors in $\mathbb{Z}_2^{\#\text{chords of } A}$ which we interpret to be characteristic vectors of the corresponding sets. The addition then corresponds to the symmetrical difference.

Now, the crossing neighborhood of y changes only if $i_{x,y} = 1$. In the case that x and y are interlaced, the crossing neighborhood of y in $A \circ x$ is precisely the symmetrical difference of N_x and N_y with two exceptions:

- $x \in N_y^{A\circ x}$ iff it is locally crossing in $A \circ x$.
  - If x is locally crossing in A, it is in N_y but not in $N_y^{A\circ x}$.
  - If x is diagonally crossing in A, it is in $N_y^{A\circ x}$ but not in N_y.
  - If x is globally crossing in A, it is neither in N_y nor in $N_y^{A\circ x}$.

  We thus have to add an additional x iff x is not globally crossing in A.
- $y \notin N_y^{A\circ x}$. N_x contains y iff y is locally crossing in A, so we have to add an additional y iff that is the case.

$g_y^{A\circ x} \equiv g_y + i_{x,y} \cdot (\ell_y + 1)$

The global matching of y changes if and only if y is interlaced with x and in that case it is interchanged with the diagonal matching. Hence, whether the global matching of y is crossing or not changes, if and only if x and y are interlaced and y is not locally crossing.

$\ell_y^{A\circ x} \equiv \ell_y + \delta_{x,y} \cdot (g_y + 1)$

The only local matching a switch at x changes is the local matching of x itself. As the local matching is interchanged with the diagonal matching, whether or not a chord y is locally crossing changes, if and only if $y = x$ and the crossing matching is not the global matching.

$c_y^{A\circ x} \equiv c_y + i_{x,y} \cdot \ell_y + \delta_{x,y} \cdot g_x$

Recall that the class is given by the local matching, or, if that is crossing, by the diagonal matching. We have three cases to analyze:

- $i_{x,y} = 1$ and $\delta_{x,y} = 0$. The global and diagonal matchings of y are interchanged. Thus, the class of y changes iff it is determined by the diagonal matching, i.e. iff y is locally crossing.
- $i_{x,y}=0$ and $\delta_{x,y}=0$. The matchings of y do not change.
- $i_{x,y}=0$ and $\delta_{x,y}=1$. The switch at x interchanges the local and diagonal matchings of x. If either was crossing, the class does not change, by definition. If the global matching is crossing, the class does change.

$i_{y,z}^{A\circ x} \equiv i_{y,z} + i_{x,y} \cdot i_{x,z} \cdot (\delta_{y,z}+1)$

Consider the following case analysis:

- $y = z$. y and z are not interlaced and this does not change as guaranteed by the $\delta_{y,z}$ in the equation.
- $x = y \neq z$. The interlacement of $x = y$ and z does not change as $i_{x,y} = 0$.
- $x \neq y \neq z \neq x$. The interlacement of y and z changes if and only if both are interlaced with x. □

The rest of the proof is just a matter of calculation: in essence we substitute all the equations from 2.25 into $P(A\circ x)$, simplify the result by applying the distributive law and canceling terms (we are over $\mathbb{Z}_2$) and apply that $P(A) \equiv 0$. This could be done by a computer algebra system but the calculation is small enough so that we can do it by hand.

The usual approach to a proof of 2.24 would be to try a case analysis. While this is possible, the number of cases to consider would be large: the proofs of the equations in 2.25 already required several case analyses, many of the them *according to different criteria*. While a proof via boolean algebra as presented below may appear of lesser aesthetic value, it is less error-prone, easier to verify and far shorter than a case analysis *with all details*. To support our claim about the length of this kind of proof, we will present the calculation in full detail, without leaving anything to the reader. This calculation concludes the proof of theorem 2.22.

Proof of 2.24. Let A be an augmented chord diagram with $P(A) \equiv 0$. We seek to calculate $P(A\circ x)$. The first step is the calculation of the parity of the common crossing neighborhood of chords y and z in $A\circ x$.

$$
\begin{aligned}
& \langle N_y^{A\circ x}, N_z^{A\circ x}\rangle \\
\equiv\ & \langle N_y + i_{x,y}\cdot(N_x + (g_x+1)\cdot x + \ell_y \cdot y), N_z + i_{x,z}\cdot(N_x + (g_x+1)\cdot x + \ell_z\cdot z)\rangle \\
\equiv\ & \langle N_y, N_z\rangle \\
& + i_{x,y}\cdot(\langle N_x, N_z\rangle + (g_x+1)\langle x, N_z\rangle + \ell_y\langle y, N_z\rangle\ ) \\
& + i_{x,z}\cdot(\langle N_x, N_y\rangle + (g_x+1)\langle x, N_y\rangle + \ell_z\langle z, N_y\rangle\ ) \\
& + i_{x,y} i_{x,z}\langle N_x + (g_x+1)\cdot x + \ell_y\cdot y, N_x + (g_x+1)\cdot x + \ell_z \cdot z\rangle
\end{aligned}
$$

We now consider the last factor separately.

$$
\begin{aligned}
& \langle N_x + (g_x+1)\cdot x + \ell_y \cdot y, N_x + (g_x+1)\cdot x + \ell_z \cdot z\rangle \\
\equiv\ & \langle N_x, N_x\rangle + (g_x+1)\langle x, N_x\rangle + \ell_y\langle y, N_x\rangle \\
& + (g_x+1)\langle N_x, x\rangle + (g_x+1)^2\langle x, x\rangle + (g_x+1)\ell_y\langle y, x\rangle \\
& + \ell_z\langle N_x, z\rangle + (g_x+1)\ell_z\langle x, z\rangle + \ell_y\ell_z\langle y, z\rangle \\
\equiv\ & \langle N_x, N_x\rangle + 0 + \ell_y i_{x,y} \\
& + 0 + (g_x+1) + (g_x+1)\ell_y\delta_{x,y} \\
& + \ell_z i_{x,z} + (g_x+1)\ell_z\delta_{x,z} + \ell_y\ell_z\delta_{y,z} \\
\equiv\ & \langle N_x, N_x\rangle + \ell_y i_{x,y} + \ell_z i_{x,z} + \ell_y\ell_z\delta_{y,z} + g_x + 1 + (g_x+1)(\ell_y\delta_{x,y} + \ell_z\delta_{x,z})
\end{aligned}
$$

Substituting, we get:

$$
\begin{aligned}
& \langle N_y^{C\circ x}, N_z^{C\circ x}\rangle \\
\equiv\ & \langle N_y, N_z\rangle \\
& + i_{x,y}\cdot(\langle N_x, N_z\rangle + (g_x+1)\langle x, N_z\rangle + \ell_y\langle y, N_z\rangle\,) \\
& + i_{x,z}\cdot(\langle N_x, N_y\rangle + (g_x+1)\langle x, N_y\rangle + \ell_z\langle z, N_y\rangle\,) \\
& + i_{x,y}i_{x,z}(\langle N_x, N_x\rangle + \ell_y i_{x,y} + \ell_z i_{x,z} + \ell_y\ell_z\delta_{y,z} + g_x + 1 + (g_x+1)(\ell_y\delta_{x,y} + \ell_z\delta_{x,z})) \\
\equiv\ & \langle N_y, N_z\rangle \\
& + i_{x,y}\cdot(\langle N_x, N_z\rangle + (g_x+1)i_{x,z}\ell_x + \ell_y i_{y,z}\,) \\
& + i_{x,z}\cdot(\langle N_x, N_y\rangle + (g_x+1)i_{x,y}\ell_x + \ell_z i_{y,z}\,) \\
& + i_{x,y}i_{x,z}(\langle N_x, N_x\rangle + \ell_y i_{x,y} + \ell_z i_{x,z} + \ell_y\ell_z\delta_{y,z} + g_x + 1) \\
& + \underbrace{i_{x,y}i_{x,z}(g_x+1)(\ell_y\delta_{x,y} + \ell_z\delta_{x,z})}_{\equiv 0 \text{ because } i_{x,y}=1 \Rightarrow \delta_{x,y}=0} \\
\equiv\ & \langle N_y, N_z\rangle + i_{x,y}\langle N_x, N_z\rangle + i_{x,z}\langle N_x, N_y\rangle \\
& + 2\cdot i_{x,y}i_{x,z}(g_x+1)\ell_x + \\
& + i_{x,y}i_{y,z}\ell_y + i_{x,z}i_{y,z}\ell_z + i_{x,y}^2 i_{x,z}\ell_y + i_{x,y}i_{x,z}^2\ell_z + \underbrace{i_{x,y}i_{x,z}\ell_y\ell_z\delta_{y,z}}_{\equiv 0 \text{ unless } y=z} \\
& + i_{x,y}i_{x,z}(\underbrace{\langle N_x, N_x\rangle + g_x}_{\equiv 0 \text{ by assumption}}) + i_{x,y}i_{x,z} \\
\equiv\ & \langle N_y, N_z\rangle + i_{x,y}\langle N_x, N_z\rangle + i_{x,z}\langle N_x, N_y\rangle \\
& + i_{x,y}i_{y,z}\ell_y + i_{x,z}i_{y,z}\ell_z + i_{x,y}i_{x,z}\ell_y + i_{x,y}i_{x,z}\ell_z + i_{x,y}\ell_y\delta_{y,z} + i_{x,y}i_{x,z}
\end{aligned}
$$

Note that in the case of $y = z$ we have

$$
\begin{aligned}
\langle N_y^{A\circ x}, N_y^{A\circ x}\rangle \ &\equiv\ \langle N_y, N_y\rangle + i_{x,y}(\ell_y+1) \\
&\equiv\ |N_y| + i_{x,y}(\ell_y+1)
\end{aligned}
$$

and we have only made use of the assumption that $\langle N_x, N_x\rangle + g_x \equiv 0$. We can now turn to the proof proper.

$$
\begin{aligned}
& P_{y,z}(A\circ x) \\
\equiv\quad & \langle N_y^{A\circ x}, N_z^{A\circ x}\rangle + i_{y,z}^{A\circ x}\cdot(c_y^{A\circ x}+c_z^{A\circ x}+1)+\delta_{y,z}\cdot g_y^{A\circ x} \\
\equiv\quad & \langle N_y, N_z\rangle + i_{x,y}\langle N_x, N_z\rangle + i_{x,z}\langle N_x, N_y\rangle \\
& + i_{x,y}i_{y,z}\ell_y + i_{x,z}i_{y,z}\ell_z + i_{x,y}i_{x,z}\ell_y + i_{x,y}i_{x,z}\ell_z + i_{x,y}\ell_y\delta_{y,z} + i_{x,y}i_{x,z} \\
& + (i_{y,z} + i_{x,y}\cdot i_{x,z}(\delta_{y,z}+1))\cdot(c_y + i_{x,y}\cdot\ell_y + \delta_{x,y}g_x + c_z + i_{x,z}\cdot\ell_z + \delta_{x,z}g_x + 1) \\
& + \delta_{y,z}\,(g_y + i_{x,y}\cdot(\ell_y+1)) \\
\equiv\quad & \langle N_y, N_z\rangle + i_{x,y}\langle N_x, N_z\rangle + i_{x,z}\langle N_x, N_y\rangle \\
& + i_{x,y}i_{y,z}\ell_y + i_{x,z}i_{y,z}\ell_z + i_{x,y}i_{x,z}\ell_y + i_{x,y}i_{x,z}\ell_z + i_{x,y}\ell_y\delta_{y,z} + i_{x,y}i_{x,z} \\
& + i_{y,z}(c_y + i_{x,y}\cdot\ell_y + \delta_{x,y}g_x + c_z + i_{x,z}\cdot\ell_z + \delta_{x,z}g_x + 1) \\
& + i_{x,y}\cdot i_{x,z}(\delta_{y,z}+1)\cdot(c_y + i_{x,y}\cdot\ell_y + \delta_{x,y}g_x + c_z + i_{x,z}\cdot\ell_z + \delta_{x,z}g_x + 1) \\
& + \delta_{y,z}\,g_y + i_{x,y}\ell_y\delta_{y,z} + i_{x,y}\delta_{y,z} \\
\equiv\quad & \underline{\langle N_y, N_z\rangle} + i_{x,y}\langle N_x, N_z\rangle + i_{x,z}\langle N_x, N_y\rangle \\
& + i_{x,y}i_{y,z}\ell_y + i_{x,z}i_{y,z}\ell_z + i_{x,y}i_{x,z}\ell_y + i_{x,y}i_{x,z}\ell_z + i_{x,y}i_{x,z} \\
& \underline{+ i_{y,z}(c_y + c_z + 1)} + i_{y,z}(i_{x,y}\cdot\ell_y + \delta_{x,y}g_x + i_{x,z}\cdot\ell_z + \delta_{x,z}g_x) \\
& + i_{x,y}\cdot i_{x,z}(\delta_{y,z}+1)\cdot(c_y + i_{x,y}\cdot\ell_y + \delta_{x,y}g_x + c_z + i_{x,z}\cdot\ell_z + \delta_{x,z}g_x + 1) \\
& \underline{+ \delta_{y,z}\,g_y} + 2\cdot i_{x,y}\ell_y\delta_{y,z} + i_{x,y}\delta_{y,z} \\
\overset{P(A)\equiv 0}{\equiv}\quad & i_{x,y}\langle N_x, N_z\rangle + i_{x,z}\langle N_x, N_y\rangle \\
& + i_{x,y}i_{y,z}\ell_y + i_{x,z}i_{y,z}\ell_z + i_{x,y}i_{x,z}\ell_y + i_{x,y}i_{x,z}\ell_z + i_{x,y}i_{x,z} \\
& + i_{y,z}(i_{x,y}\cdot\ell_y + \delta_{x,y}g_x + i_{x,z}\cdot\ell_z + \delta_{x,z}g_x) \\
& + i_{x,y}\cdot i_{x,z}(\delta_{y,z}+1)\cdot(c_y + i_{x,y}\cdot\ell_y + \delta_{x,y}g_x + c_z + i_{x,z}\cdot\ell_z + \delta_{x,z}g_x + 1) \\
& + i_{x,y}\delta_{y,z} \\
\equiv\quad & i_{x,y}\langle N_x, N_z\rangle + i_{x,z}\langle N_x, N_y\rangle \\
& + i_{x,y}i_{y,z}\ell_y + i_{x,z}i_{y,z}\ell_z + i_{x,y}i_{x,z}\ell_y + i_{x,y}i_{x,z}\ell_z + i_{x,y}i_{x,z} \\
& + i_{y,z}\cdot i_{x,y}\cdot\ell_y + i_{y,z}\cdot i_{x,z}\cdot\ell_z + i_{y,z}(\delta_{x,y}g_x + \delta_{x,z}g_x) + \underbrace{i_{x,y}\cdot i_{x,z}(\delta_{x,y}g_x + \delta_{x,z}g_x)}_{\equiv 0} \\
& + i_{x,y}\cdot i_{x,z}\cdot(c_y + c_z + 1) + i_{x,y}\cdot i_{x,z}\cdot(i_{x,y}\cdot\ell_y + i_{x,z}\cdot\ell_z) \\
& + i_{x,y}\cdot i_{x,z}\cdot\delta_{y,z}\cdot\underbrace{(c_y + i_{x,y}\cdot\ell_y + \delta_{x,y}g_x + c_z + i_{x,z}\cdot\ell_z + \delta_{x,z}g_x + 1)}_{\equiv 1 \text{ if } y=z} \\
& + i_{x,y}\delta_{y,z} \\
\equiv\quad & i_{x,y}\langle N_x, N_z\rangle + i_{x,z}\langle N_x, N_y\rangle \\
& + i_{x,y}i_{y,z}\ell_y + i_{x,z}i_{y,z}\ell_z + i_{x,y}i_{x,z}\ell_y + i_{x,y}i_{x,z}\ell_z + i_{x,y}i_{x,z} \\
& + i_{x,y}i_{y,z}\ell_y + i_{x,z}i_{y,z}\ell_z + i_{y,z}(\delta_{x,y}g_x + \delta_{x,z}g_x) \\
& + i_{x,y}\cdot i_{x,z}\cdot(c_y + c_z + 1) + i_{x,y}i_{x,z}\ell_y + i_{x,y}i_{x,z}\ell_z \\
& + \underbrace{i_{x,y}i_{x,z}\delta_{y,z} + i_{x,y}\delta_{y,z}}_{\equiv 0}
\end{aligned}
$$

$$
\begin{aligned}
&\equiv && i_{x,y}\langle N_x, N_z\rangle + i_{x,z}\langle N_x, N_y\rangle \\
& && + i_{x,y}i_{x,z} + i_{y,z}(\delta_{x,y}g_x + \delta_{x,z}g_x) + i_{x,y}\cdot i_{x,z}\cdot(c_y + c_z + 1) \\
& && \ell_y(2\cdot i_{x,y}i_{y,z} + 2\cdot i_{x,y}i_{x,z}) + \ell_z(2\cdot i_{x,z}i_{y,z} + 2\cdot i_{x,y}i_{x,z}) \\
&\equiv && i_{x,y}\langle N_x, N_z\rangle + i_{x,z}\langle N_x, N_y\rangle \\
& && + i_{x,y}i_{x,z} + i_{y,z}(\delta_{x,y}g_x + \delta_{x,z}g_x) + i_{x,y}\cdot i_{x,z}\cdot(c_y + c_z + 1) \\
&\overset{*}{\equiv} && i_{x,y}\cdot(i_{x,z}\cdot(c_x + c_z + 1) + \delta_{x,z}\cdot g_x) + i_{x,z}\cdot(i_{x,y}\cdot(c_x + c_y + 1) + \delta_{x,y}\cdot g_x) \\
& && + i_{x,y}i_{x,z} + i_{y,z}(\delta_{x,y}g_x + \delta_{x,z}g_x) + i_{x,y}\cdot i_{x,z}\cdot(c_y + c_z + 1) \\
&\equiv && i_{x,y}\cdot i_{x,z}\cdot(2c_x + 2c_z + 2c_y + 3) + i_{x,y}i_{x,z} \\
& && + i_{x,y}\cdot\delta_{x,z}\cdot g_x + i_{x,z}\cdot\delta_{x,y}\cdot g_x \\
& && + i_{y,z}(\delta_{x,y}g_x + \delta_{x,z}g_x) \\
&\equiv && g_x\cdot(\delta_{x,z}\cdot i_{x,y} + \delta_{x,y}\cdot i_{x,z} + \delta_{x,y}\cdot i_{y,z} + \delta_{x,z}\cdot i_{y,z}) \\
&\overset{**}{\equiv} && g_x\cdot(\delta_{x,z}\cdot i_{x,y} + \delta_{x,y}\cdot i_{x,z} + \delta_{x,y}\cdot i_{x,z} + \delta_{x,z}\cdot i_{x,y}) \\
&\equiv && 0
\end{aligned}
$$

where at $*$ we make use of the equations

$$
\begin{aligned}
\langle N_x, N_z\rangle &\equiv i_{x,z}\cdot(c_x + c_z + 1) + \delta_{x,z}\cdot g_x \\
\langle N_x, N_y\rangle &\equiv i_{x,y}\cdot(c_x + c_y + 1) + \delta_{x,y}\cdot g_x
\end{aligned}
$$

which hold by the assumption that $P(A)\equiv 0$. At $**$ we use of the fact that $x = y \Rightarrow i_{y,z} = i_{x,z}$, and $x = z \Rightarrow i_{y,z} = i_{y,x} = i_{x,y}$.

Again we make a note about the special case $y = z$. In this case we could calculate

$$
\begin{aligned}
& \langle N_y^{A\circ x}, N_y^{A\circ x}\rangle + g_y^{A\circ x} \\
\equiv\ & \langle N_y, N_y\rangle + i_{x,y}(\ell_y + 1) + g_y + i_{x,y}\cdot(\ell_y + 1) \\
\equiv\ & \langle N_y, N_y\rangle + g_y \\
\equiv\ & 0
\end{aligned}
$$

in which case we have only made use of the assumptions $\langle N_y, N_y\rangle + g_y \equiv \langle N_x, N_x\rangle + g_x \equiv 0$. So we can remark that the conjunction of the equations

$$\langle N_x, N_x\rangle + g_x \equiv 0$$

for all chords x forms an invariant under the switch operation in itself. □

Building upon the work of de Fraysseix and Ossona de Mendez, we have thus developed a powerful invariant under the switch operation, using which we were able to prove a Rosenstiehl-type criterion for the realizability of augmented chord diagrams, which in turn implies the Rosenstiehl criterion for cross-realizability of chord diagrams. Using the “smaller” invariant mentioned above, we can complete our proof of 2.13. We first show the following statement about augmented chord diagrams:

2.26. Let A be an augmented chord diagram and v_1, ..., v_k its locally crossing chords. A is realizable if and only if the following two conditions hold:

i. no two chords of $A \circ v_1 \ldots \circ v_k$ with the same local color are interlaced

ii. every chord x the crossing neighborhood of x in A is odd iff x is globally crossing in A

Proof. We already know that the specified conditions are necessary. All we have to show is that they are sufficient. Let A be an augmented chord diagram for which both conditions hold. By 2.20 all we have to show is that no chord is globally crossing in $A \circ v_1 \ldots \circ v_k$. As we have seen the conjunction of the equations

$$\langle N_x, N_x \rangle + g_x \equiv 0$$

is an invariant of the switch operation, and this conjunction is precisely the statement that the crossing neighborhood of every chord x is odd iff x is globally crossing. So this statement holds in $A \circ v_1 \ldots \circ v_k$ iff it holds in A, which it does by assumption. As the crossing neighborhood of every chord in $A \circ v_1 \ldots \circ v_k$ is empty, it follows that no chord is globally crossing as required. □

We can now prove 2.13, i.e. a chord diagram C with chords v_1, ..., v_k is cross-realizable if and only if C is even and $C \circ v_1 \ldots \circ v_k$ is bipartite.

Proof of 2.13. Again, we already know that 2.13 is necessary and we have to show that it is also sufficient. By definition, there exists a crossing augmentation A of C that is even, if and only if C is even. By 2.21, there exists a crossing augmentation A of C with the property that no two chords of $A \circ v_1 \ldots \circ v_k$ with the same local color are interlaced, if and only if $C \circ v_1 \ldots \circ v_k$ is bipartite. Now we observe that for crossing augmentations A 2.26.ii is fulfilled if and only if A is even. Taking these statements together, we obtain that there exists a crossing augmentation A of C that meets the conditions of 2.26 if and only if C meets the conditions of 2.13 and so 2.13 follows from 2.26. □

2.7 The Lovász-Marx Criterion for Cross-Realizability

In the previous section we have shown a combinatorial criterion for the realizability of augmented chord diagrams which implies the Rosenstiehl Criterion for the cross-realizability of chord diagrams (and interlacement graphs). In this section we will obtain a proof of the Lovász-Marx Criterion for cross-realizability of interlacement graphs from the Rosenstiehl Criterion. The proof we present is due to Aigner. It is the only proof of the Lovász-Marx Criterion the author was able to find in the literature, as Lovász and Marx do not give a proof of their theorem in [6].

We have seen characterizations of realizable chord diagrams and realizable augmented chord diagrams. This time we want to focus on interlacement graphs. Recall that we defined a partial order on the set of chord diagrams via two operations, the loop removal and the switch deletion, that each reduce the number of chords in the given diagram. A chord x splits a curve into two closed segments c_1 and c_2. The loop removal at a chord x required us to remove x along with all chords incident with one of the two segments, say c_1.

Different chord diagrams may have the same interlacement graph Λ. Moreover, the vertices $y \neq x$ of Λ may belong to c_1, to c_2, or to both and for a given vertex y it is not uniquely determined by Λ which one of these cases applies. So the loop removal does not induce a well defined operation on interlacement graphs. We will therefore consider a modified concept of loop removal: instead of deleting x and all chords incident to one segment, we remove x and all chords interlaced with x. For Λ this means that we delete the vertex x along with its neighborhood.

Does this modified version also preserve cross-realizability? Geometrically, the new concept of loop removal corresponds to shrinking c_1 until it is contained entirely in one face of c_2. This removes precisely the crossings between the two global segments. We then perform a switch deletion at v which only has the effect of removing v since at this point v does not have any neighbors. An example is given in Figure 2.28. We conclude that the modified concept of loop removal also preserves cross-realizability. For interlacement graphs Λ and Λ', we write $\Lambda' \leqslant_r \Lambda$ if Λ' can be obtained from Λ by a (possibly empty) sequence of loop removals.

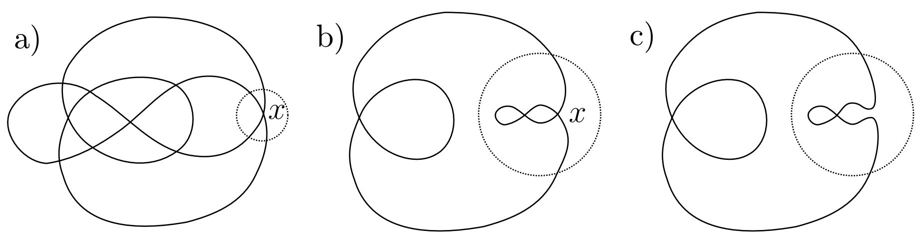

Figure 2.28. The modified loop removal operation at x works like this: we start out with a curve a), shrink one of the global segments at x to obtain b) and then we apply the switch removal to get the result c).

In [6] Lovász-Marx formulate their theorem on the level of chord diagrams using the original version of loop removal. We are going to prove an analogue of their theorem due to Aigner [1] that is formulated on the level of interlacement graphs and uses the modified concept of loop removal.

The switch deletion on chord diagrams does induce induce a well defined operation on interlacement graphs: switch deleting the vertex x means replacing every edge with a non-edge (and vice versa) in the neighborhood of x and then deleting x. We write $\Lambda' \leqslant_s \Lambda$ if Λ' can be obtained from Λ by a (possibly empty) sequence of switch deletions.

For the purposes of this section we define the relation $\leqslant$ on the set of all graphs, as the order relation generated by the union of $\leqslant_r$ and $\leqslant_s$. In other words, $\Lambda' \leqslant \Lambda$ iff Λ' can be obtained from Λ by a sequence of switch deletions and loop removals. $\leqslant$ is anti-symmetric because both operations strictly reduce the number of vertices. The set of interlacement graphs (also called circle graphs) is the set of all $\mathcal{I}(C)$ for some chord diagram C and it is denoted by $\mathcal{C}i$. By construction we have:

2.27. Let $\Lambda', \Lambda \in \mathcal{C}i$ be interlacement graphs with $\Lambda' \leqslant \Lambda$. If Λ is cross-realizable, so is Λ'.

The variant of the Lovász-Marx Criterion we are going to show now is:

2.28. Lovász-Marx Criterion for Interlacement Graphs AIGNER

$\Lambda \in \mathcal{C}i$ is cross-realizable iff $K_{2n} \nleqslant \Lambda$ for all $n \geqslant 1$.

Again, that the condition is necessary follows immediately from 2.27 and the fact that no K_{2n} is cross-realizable. To show that the condition is sufficient we introduce further classes of graphs. Let $\mathcal{K}$ be the class of graphs G with $K_{2n} \nleqslant G$ for all $n \geqslant 1$. Let $\mathcal{E}$ be the class of graphs in which every vertex is even. Finally, let $\mathcal{R}$ be the class of graphs for which the Rosenstiehl condition holds, i.e. all graphs G such that a) every vertex is even, b) every non-edge is even and c) the set of even edges forms a cut in G. Building on the fact that a graph is bipartite if and only if it contains no odd circle, we note that c) is equivalent to the statement: G contains no circle with an odd number of even edges.

Given this notation, we can immediately observe the following:

1. By the Rosenstiehl Criterion for cross-realizability, the class of cross-realizable interlacement graphs is just $\mathcal{C}i \cap \mathcal{R}$.

2. Every Rosenstiehl graph is even by definition, i.e. $\mathcal{R} \subset \mathcal{E}$.

3. Theorem 2.28 is equivalent to $\mathcal{C}i \cap \mathcal{R} = \mathcal{C}i \cap \mathcal{K}$.

4. If Λ is realizable, then $K_{2n} \nleqslant \Lambda$. Hence $\mathcal{C}i \cap \mathcal{R} \subset \mathcal{K}$ and all that is left to show is $\mathcal{C}i \cap \mathcal{K} \subset \mathcal{C}i \cap \mathcal{R}$.

We now come to the first of two lemmas that constitute the proof of 2.28.

2.29. $\mathcal{E}$ is $\leqslant_s$-closed, and $G \in \mathcal{E} \Longleftrightarrow K_{2n} \nleqslant_s G$ for all $n \geqslant 1$.

Proof. *$\mathcal{E}$ is $\leqslant_s$-closed.* Let $G \in \mathcal{E}$ and H be the graph obtained from G by switch deletion of u. We have to show that all vertices of H have even degree, which is trivially satisfied for all vertices not adjacent to u in G. Let v be a vertex adjacent to u. The set of neighbors of v in H is the symmetrical difference of $N_G(v)$ and $N_G(u)$, without v and u. So we have

$$\begin{aligned} |N_H(v)| &= |N_G(v)| + |N_G(u)| - 2|N_G(v) \cap N_G(u)| - 2 \\ &\equiv |N_G(v)| + |N_G(u)| \equiv 0 \pmod 2 \end{aligned} \tag{2.6}$$

$G \in \mathcal{E} \Longrightarrow K_{2n} \nleqslant_s G$ for all $n \geqslant 1$. As no K_{2n} is even, this follows from the fact that $\mathcal{E}$ is $\leqslant_s$-closed.

$G \notin \mathcal{E} \Longrightarrow K_{2n} \leqslant_s G$ for some $n \geqslant 1$. It suffices to show the implication for all $\leqslant_s$-minimal graphs with $G \notin \mathcal{E}$. Let G be such a graph, which is connected by $\leqslant_s$-minimality. As G is not even, there is a vertex u of odd degree. Let v be one of its neighbors. The switch deletion at u yields a graph H, which is even by $\leqslant_s$-minimality of G. We calculate

$$0 \equiv |N_H(v)| \overset{(2.6)}{\equiv} |N_G(v)| + |N_G(u)| \equiv |N_G(v)| + 1 \pmod 2$$

and conclude that v is odd in G. As G is connected it follows that all vertices of G are odd. Assume there are vertices x and y in G that are not adjacent. The switch deletion at x produces a graph H' and does not affect the degree of y, so we have $0 \equiv |N_{H'}(y)| \equiv |N_G(y)| \equiv 1$, a contradiction. It follows that G is complete, all vertices are odd and hence $G = K_{2n}$ for some $n \geqslant 1$. □

Apart from this important observation about even graphs, the key ingredient in our proof is the following: we already know that $\leqslant$ preserves cross-realizability. Of course $\leqslant$ does not preserve non-cross-realizability, as K_{2n} is not cross-realizability, but all $G < K_{2n}$ are cross-realizability. So the fact that all $<$-minors of a graph G are cross-realizable tells us nothing about the cross-realizability of G. To understand the Lovász-Marx theorem, we need to understand in what situations we *can* say something about the cross-realizability of G. It turns out that these situations are surprisingly easy to characterize.

2.30. Let G be a graph, such that $H < G$ implies $H \in \mathcal{R}$. Then $G \in \mathcal{R} \Longleftrightarrow G \in \mathcal{E}$.

Note that if we restrict our attention to interlacement graphs G, then 2.30 is indeed a statement about the cross-realizability of G instead of about whether or not the Rosenstiehl condition holds for G.

Proof. If $G \in \mathcal{R}$ then of course $G \in \mathcal{E}$. So let G be a graph with $\forall H\colon H < G \Rightarrow H \in \mathcal{R}$ and $G \in \mathcal{E}$, i.e. Rosenstiehl condition a) is satisfied by assumption.

Rosenstiehl condition b) holds for G. Let u, v be two non-adjacent vertices of G. The loop removal at u decreases the degree of v by the size of their common neighborhood. If H denotes the graph thus obtained, we have

$$\begin{aligned} & |N_H(v)| = |N_G(v)| - |N_G(v) \cap N_G(u)| \\ \Rightarrow \quad & |N_G(v) \cap N_G(u)| \equiv |N_G(v)| + |N_H(v)| \equiv 0 \pmod 2 \end{aligned}$$

which means that u and v have an even number of common neighbors as desired. Note that all congruences in this proof are modulo 2.

Rosenstiehl condition c) holds for G. We want to check that every cycle in G has an even number of even edges. As the cycle space is generated by the induced cycles, we can restrict our attention to those. To be able to count the number of even edges we introduce a function $\gamma(u, v): V \times V \rightarrow \mathbb{N}$ defined by

$$\gamma(u, v) = |N_G(u) \cap N_G(v)| + 1 \tag{2.7}$$

so that the pair (u, v) is even iff $\gamma(u, v) \equiv 1 \pmod 2$. Consequently, a cycle $Z = v_1 ... v_k$ (where all v_i are different and v_1 and v_k adjacent) has even many even edges if and only if

$$\Sigma(Z) := \sum_{i=1}^{k} \gamma(v_i, v_{i+1}) \equiv 0$$

where $v_{k+1} := v_1$. Furthermore we define for a fixed set $Y = \{a, b, c\}$ of 3 vertices and for any subset $X \subset Y$ the quantity $t_Y(X)$ to be the number of vertices $v \notin Y$ that are adjacent to all $x \in X$ but to no other vertex $y \in Y \setminus X$. Given this notation we have, for example, for two vertices $a, b \in Y$

$$|N_G(a) \cap N_G(b)| = t(a, b) + t(a, b, c) + \begin{cases} 1, \text{ if } c \text{ is adjacent to both } a \text{ and } b \\ 0, \text{ otherwise} \end{cases} \tag{2.8}$$

Now, let Z be any induced cycle in G. We will show that $\Sigma(Z) \equiv 0$ by considering two cases.

Case 1: $Z = vuw$ is a triangle. In this case $t := t_{vuw}$. As any two vertices have the third as a common neighbor, we have by 2.7 and 2.8:

$$\begin{aligned} \gamma(u, v) &= t(u, v) + t(u, v, w) + 2 \\ \gamma(u, w) &= t(u, w) + t(u, v, w) + 2 \\ \gamma(v, w) &= t(v, w) + t(u, v, w) + 2 \end{aligned}$$

$$\Rightarrow \quad \Sigma(Z) \equiv t(u, v) + t(u, w) + t(v, w) + t(u, v, w)$$

Since G is even, we have

$$\begin{aligned} 0 &\equiv |N_G(u)| = t(u) + t(u,v) + t(u,w) + t(u,v,w) + 2 \\ \Rightarrow \quad \Sigma(Z) &\equiv t(u) + t(v,w) \end{aligned}$$

Let H denote the graph obtained from G by switch deletion at u. We now consider the common neighborhood of v and w in H. It consists of the common neighbors of v and w in G that were not adjacent to u, and the neighbors of u that were not adjacent to either v or w in G. As v and w are not adjacent in $H \in \mathcal{R}$, their common neighborhood is even and so

$$\begin{aligned} \Sigma(Z) &\equiv t(u) + t(v,w) \\ &\equiv |N_H(v) \cap N_H(w)| \\ &\equiv 0 \end{aligned}$$

Case 2: Z has length $\geqslant 4$. This means $Z = vuwx_1...x_k$ for some $k \geqslant 1$. Again, we consider the graph H obtained by switch deletion at u which contains the cycle $Z' = vwx_1...x_k$. Our aim is to show $\Sigma_G(Z) \equiv \Sigma_H(Z')$. We first note that the switch deletion of u does not affect the number of common neighbors of x_i and x_{i+1} as u is adjacent to neither, so

$$\begin{aligned} \Sigma_G(Z) + \Sigma_H(Z') \equiv\ & \gamma(x_k,v) + \gamma(v,u) + \gamma(u,w) + \gamma(w,x_1) \\ & + \gamma_H(x_k,v) + \gamma_H(v,w) + \gamma_H(w,x_1) \end{aligned} \tag{2.9}$$

As in the previous case we have

$$\begin{aligned} \gamma(u,v) &= t_{vuw}(u,v) + t_{vuw}(u,v,w) + 1 \\ \gamma(u,w) &= t_{vuw}(u,w) + t_{vuw}(u,v,w) + 1 \\ \gamma_H(v,w) &= t_{vuw}(u) + t_{vuw}(v,w) + 1 \end{aligned}$$

where t counts neighbors in G. As G satisfies Rosenstiehl conditions a) and b)

$$\begin{aligned} 0 \equiv |N_G(u)| &\equiv t_{vuw}(u) + t_{vuw}(u,v) + t_{vuw}(u,w) + t_{vuw}(u,v,w) + 2 \\ 0 \equiv |N_G(v) \cap N_G(w)| &\equiv t_{vuw}(v,w) + t_{vuw}(u,v,w) + 1 \end{aligned}$$

which, when summing all 5 congruences, gives

$$\gamma_H(v,w) \equiv \gamma(v,u) + \gamma(u,w)$$

Now we proceed by considering the three vertices u, w, x_1. As x_1 is not adjacent to u, the common neighborhood of w and x_1 in H is the set of vertices a with $a \sim w$ and $a \sim x_1$ but $a \nsim u$ in G, plus the set of vertices b with $b \nsim w, b \sim u$ and $b \sim x_1$, as these are not removed from the neighborhood of w and remain in the neighborhood of x_1.

$$\begin{aligned} \gamma(w, x_1) &= t_{uwx_1}(w, x_1) + t_{uwx_1}(u, w, x_1) + 1 \\ \gamma_H(w, x_1) &= t_{uwx_1}(w, x_1) + t_{uwx_1}(u, x_1) + 1 \end{aligned}$$

Since $ux_1 \notin E(G)$ and b) holds for G we also have

$$\begin{aligned} 0 \equiv |N_G(u) \cap N_G(x_1)| &\equiv t_{uwx_1}(u, x_1) + t_{uwx_1}(u, w, x_1) + 1 \\ \Rightarrow \qquad \gamma_H(w, x_1) &\equiv \gamma(w, x_1) + 1 \end{aligned}$$

By symmetry, the congruence $\gamma_H(x_k,\ v) \equiv \gamma(x_k,\ v) + 1$ holds as well. We now have expressed $\gamma_H(x_k,\ v)$, $\gamma_H(v,\ w)$ and $\gamma_H(w,\ x_1)$ in terms of γ. Substituting these results in (2.9) we obtain the desired result $\Sigma_G(Z) + \Sigma_{G\circledcirc u}(Z') \equiv 0$. □

It is now easy to complete the proof of our version of the Lovász-Marx Criterion.

Proof of 2.28. We show $\mathcal{K} \subseteq \mathcal{R}$, which means that if $K_{2n} \nleqslant G$ for all $n \geqslant 1$, then $G \in \mathcal{R}$. Assume the contrary is true and pick a $\leqslant$-minimal counterexample G, i.e. $K_{2n} \nleqslant G$ for all $n \geqslant 1$, $G \notin \mathcal{R}$, but $H \in \mathcal{R}$ for all $H < G$. From $K_{2n} \nleqslant G$ follows $K_{2n} \nleqslant_s G$ for all $n \geqslant 1$, which implies $G \in \mathcal{E}$ by 2.29. But then we can apply 2.30 to obtain $G \in \mathcal{R}$, which is a contradiction. □

2.8 The Number of Realizations

2.31. A realizable augmented chord diagram (C, a) has exactly one realization c.

Proof of 2.31. First of all we note that we can ignore the traversal of the graph $G = \mathcal{G}(C)$ given by C, as it must be the same for all realizations of C. We are only interested in showing that the drawing of G in the plane is uniquely determined by a. So instead of considering $(C,\ \pi)$ and $(C,\ a)$ we can consider $(G,\ \pi)$ and $(G,\ a)$. We then have to show that given a realizable augmentation a of G, the rotation system π of a realization is uniquely determined by a. By 1.11 it then follows that the equivalence class of a realization (G, d) is uniquely determined.

Therefore, let (G, d) be a realization of (G, a) in the plane with rotation system $\pi = (\pi_v)_v$. The augmentation a defines a proper coloring of the faces of (G, d) with black and white. At any vertex v the edges of the two touching matchings form a cycle on the half-edges incident with v (see Figure 2.29). This cycle gives the undirected cyclic order of the half edges in the local rotation at v. This means that at every vertex v the local rotation π_v is determined by a_v up to reversal.

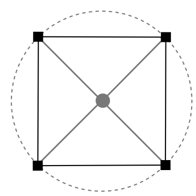

Figure 2.29. The edges of a touching matching at a double point v form a cycle on the set of half-edges incident with v.

Now, pick a vertex v_0. Fix an orientation of the local rotation π_{v_0} at v_0. If we can now show that the orientation of all other local rotations π_v are uniquely determined by π_{v_0} we are done. To see this, we argue that if vertices v and w are adjacent, then a_v, a_w and the oriented local rotation at v determine the oriented local rotation at w:

Consider Figure 2.30. Denote the half-edges of the edge vw with h_v and h_w. Denote the edges of the black matching that are incident with h_v and h_w with b_{h_v} and b_{h_w}, respectively. b_{h_v} and b_{h_w} belong to the same face of (G, d), because vw bounds on two faces with different color, and hence b_{h_v} and b_{h_w} are on the "same side" of vw. This means that if b_{h_v} is oriented towards h_v by π_v, b_{h_w} has to be oriented away from h_w by π_w, as π is the rotation system of (G, d). Correspondingly, if b_{h_v} is oriented away from h_v by π_v, b_{h_w} has to be oriented towards h_w by π_w. Thus the orientation of π_w is determined by π_v.

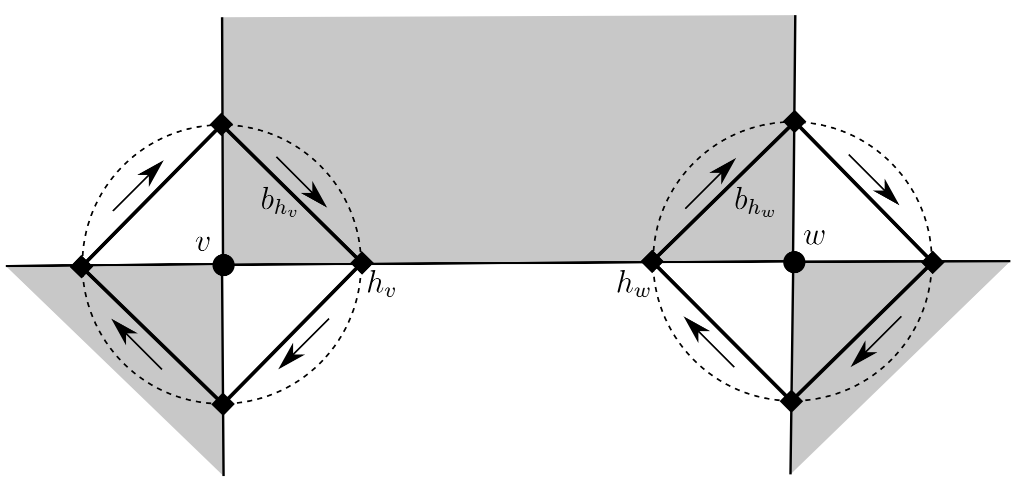

Figure 2.30. Given an augmentation $(a_u)_u$ and a local rotation π_v at v, the local rotation π_w at w is uniquely determined.

Note that the same does not hold if (G, a) is not realizable. While (G, a) always defines a local rotation at every vertex up to reversal, the relation expressed in Figure 2.30 does not in general lead to a consistent orientation of the local rotations. □

Another way to put 2.31 is to say that the function mapping a curve c to its augmented chord diagram $\mathcal{A}(c)$ has an inverse (which is given by the proof of 2.31) and, hence, is bijective. We can use this fact to count the number of realizations of a chord diagram C with a given type function ϑ.

2.32. Let C be a chord diagram and ϑ a type vector. Then

$$\#\text{curves realizing } (C, \vartheta) = \#\text{realizable augmentations of } (C, \vartheta)$$

Proof. By 2.31 the set of curves realizing $(C,\ \vartheta)$ and the set of realizable augmentations of (C, ϑ) are in bijection. □

In 1.14 we computed the number of cross-realizations of a cross-realizable chord diagram C to be 2^{k-1}, where k is the number of components of $\mathcal{I}(C)$. There we used a geometric argument to show that there are indeed that many non-equivalent cross-realizations. Using 2.32 and the equational version 2.23 of our criterion for the realizability of augmented chord diagrams, we can now easily compute that the same formula holds for the number of ϑ-realizations of C for a given type vector ϑ.

2.33. Let C be a chord diagram that is ϑ-realizable for some type vector ϑ. Denote the number of components of $\mathcal{I}(C)$ with k. Then

$$\#\text{curves realizing } (C, \vartheta) \ = \ 2^{k-1}$$

Proof. By 2.32 and the equational version of our criterion for the realizability of augmented chord diagrams 2.23, we need to compute the number of augmentations $A = (C, (a_v)_v)$ of (C, ϑ), with the property that the equation

$$\langle N_x, N_y\rangle + i_{x,y} \cdot (c_x + c_y + 1) \equiv 0$$

holds for all pairs of chords $x \neq y$ and the equation

$$\langle N_x, N_x\rangle + g_x \equiv 0$$

holds for all chords x.

Note that the size of the common crossing neighborhood $\langle N_x, N_y\rangle$ is determined by C and ϑ. So for a given x the value of g_x is the same for all realizable augmentations A. As $i_{x,y}$ is also determined by C, the only respect in which the realizable augmentations A can differ is with regard to the choice of the c_x, the classes of the individual chords. If two chords are interlaced, we have

$$c_x + c_y \equiv \langle N_x, N_y\rangle + 1$$

so the class c_x of a chord x uniquely determines the classes of all chords y in the same component as x. We thus can have at most 2^k such augmentations A. However, an augmentation $(a_v)_v$ is identified with the augmentation $(s \circ a_v)_v$ where s is the bijection s: {black, white, crossing} $\to$ {black, white, crossing} that swaps black and white. The effect of s on our variables is that the value of each c_x is flipped, i.e. $s: c_x \mapsto c_x + 1 \pmod 2$ for all x. Thus, we have counted every equivalence class twice and, hence, obtain an upper bound of 2^{k-1}.

All we have to show now is that there are indeed 2^{k-1} different augmentations with the above properties. To that end we fix a chord x and an augmentation A_0 realizing (C, ϑ). We use A_0 to denote the representative of the equivalence class of augmentations $[A_0]$ with $c_x = 0$. Denote the k components of $\mathcal{I}(C)$ with C_1, ..., C_k such that $x \in C_1$. For every vector $v \in \{0, 1\}^{k-1}$ we define $A + v$ to be the augmentation obtained from A by mapping $c_y \mapsto c_y + v_i$ for all chords $y \notin C_1$ where i is the index of the component containing y. All $A_0 + v$ are augmentations of (C, ϑ) as only the classes of the chords changed. We claim that all $A_0 + v$ are realizable. For every individual chord z the equation holds, because it holds for A_0, and for every pair of chords $z_1 \neq z_2$ that are in the same component C_i it holds by construction. Now, let $z_1 \neq z_2$ be chords in two different components. As they are in different components they are not interlaced and hence the corresponding equation has the form

$$\langle N_x, N_y \rangle \equiv 0$$

which holds, because it holds for A_0. Hence all equations hold and by 2.23 all $A_0 + v$ are realizable. For $v \neq v'$ the two augmented chord diagrams $A_0 + v$ and $A_0 + v'$ are indeed in different equivalence classes, because they differ with regard to the class c_z of some chord z, but agree with regard to the class c_x of the fixed chord x. This proves the formula as $|\{0, 1\}^{k-1}| = 2^{k-1}$. □

2.9 Outlook: Sets of Curves

In this chapter we saw numerous applications of the switch operation which is a useful concept for the study of curves in the plane. We explicitly considered curves that can have both touching and/or crossing double points and defined the switch operation on this larger domain. By the introduction of augmented chord diagrams we were able to keep track of the type of the individual double points, with the consequence that the realizability of an augmented chord diagram is invariant under the switch operation. This is a very nice property that, of course, substructure operations such as the switch deletion do not have.

The switch corresponds to replacing the local matching at a double point v with the diagonal matching. As we already observed in section 2.2, the replacement of the local matching at v with the global matching at v yields two curves, instead of one. Therefore, we cannot apply this operation as it leads out of the domain of objects we consider. However, in terms of substructure operations, this corresponds to the loop removal which we met in the context of the Lovász-Marx Criterion and which is a very useful concept as well. So, as with the switch, we would like to come up with an analogue of the loop removal under which realizability is invariant.

In case of the switch, we avoided the deletion of the touching double point that arises from the switch at a crossing double point by extending the domain we consider to curves that may have double points. In case of the loop removal, we now would like to avoid the deletion of one of the loops by considering *sets of curves* as opposed to individual curves. We now present a brief outlook on this subject matter.

Note, however, that our purpose here is not to use the "switch-approach" to give criteria for the realizability of more general objects than curves. To that end, we refer the reader to Fleming and Mellor [4] who employ the "switch-approach" to give an "algorithmic" criterion for the cross-realizability of directed graphs in general. We rather argue that a setting that allows for a combination of the approaches taken by de Fraysseix-Ossona de Mendez on the one hand and Lovász-Marx on the other is beneficial for the study of curves themselves.

An **Euler partition** of a multigraph G is a set of closed walks $w = \{w_i\colon 1 \leqslant i \leqslant k\}$ on G with the property that every half-edge of G is in exactly one of the walks w_i. A **set of curves** in the plane is a triple (G, d, w) where G is a connected multigraph, d is a cellular embedding of G in the plane and w is an Euler partition of G. Let M_v denote the set of matchings at a vertex v. Given all of our previous observations it is immediate that a function $\ell\colon V(G) \to \bigcup_{v\in V(G)} M_v$ with $\ell(v) \in M_v$ that selects a local matching at each of the vertices v defines an Euler partition of G and vice versa. We call the operation induced on sets of curves by changing the value of $\ell(v)$ at a single vertex a **local replacement**.

Note that the terms local and global matching are still well-defined in this context. However, at a double point where two different curves meet the local and the global matching coincide (we call such a double point **disconnecting**). Thus, the set of walks does not define a 3-coloring of the matchings at a double point. The drawing d, however, still does and hence the concept of an augmentation that colors the matchings with "crossing", "black" and "white" is well-defined.

A **generalized chord diagram** is a pair of multigraphs (G_r, G_c) on the same vertex set such that G_r is 2-regular and G_c is 1-regular. This is just the definition of a chord diagram, except that the requirement that G_r is a circle has been weakened. With generalized chord diagrams the components of G_r are circles and every circle corresponds to one walk in w. So chord diagrams are generalized chord diagrams that represent a set of a single curve. Just as in the case of chord diagrams, we have a one-to-one correspondence between generalized chord diagrams and pairs $(G,\ w)$ of 4-regular multigraphs G and Euler partitions w. Note that below we consider only the case that G is connected.

The augmentation of a set of curves depends only on the pair (G, d). If we now define **augmented generalized chord diagrams**, the induced local replacement operation on these objects will automatically leave realizability invariant as the local replacement only affects the set of walks w and not the augmentation.

2.34. Any augmented generalized chord diagram A can be converted into a touching augmented chord diagram A' by a sequence of local replacements.

A is realizable if and only if no two chords of the same color are interlaced in A' and no chord is globally crossing.

Proof. At any chord v of A there is always one matching that is neither global nor crossing. We now replace the local matching of v with that matching and call the resulting diagram A_1. Note that if any chord w is disconnecting in A_1, it was also disconnecting in A. Iterating this process, we obtain a touching augmented chord diagram A'. A is realizable if and only if A' is realizable because realizability is invariant under the local replacement operation. By 2.19 the result follows. □

This, again is an "algorithmic" criterion for realizability. What we would like to have is a combinatorial, Rosenstiehl-type criterion. However, this is not as straightforward to obtain as one might hope. The key ingredient in the proof of 2.23 was that we were able to express the values of the boolean predicates in $A \circ v$ in terms of the values of these predicates in A (see 2.25). In this more general setting, the author was not able to define a set of predicates that permit a lemma such as 2.25.

Chapter 3
Thrackles

3.1 Definition

A **thrackle** is a drawing of a graph in the plane in which every two edges have precisely one point in common:

- if the two are incident, they share an end-point,
- otherwise, they share an interior point at which they cross each other.

We call this requirement the **thrackle condition**. A graph is **thrackleable** if it has a thrackle drawing (i.e. if it can be drawn as a thrackle). In most cases we will denote a thrackle drawing by a capital letter T to distinguish them from plane drawings d. It is immediate from the definition that deleting (the drawings of) edges from a thrackle always yields a thrackle drawing of the corresponding subgraph and therefore

---

3.1. Thrackleability is closed under the subgraph relation.

Let us now consider some examples.

The 3-circle C_3 is thrackleable, as every two of its three edges are incident. Thus, any plane drawing of the C_3 gives us a thrackle drawing as no additional crossing is required. The C_5 is a more interesting example of a thrackleable graph – a thrackle drawing is given in Figure 3.1.

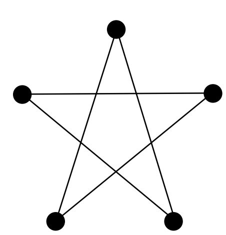

Figure 3.1. A thrackle drawing of C_5.

As it turns out, the $C_4 = (\{a,b,c,d\}, \{ab,bc,cd,da\})$ is not thrackleable (see Figure 3.2a). To see this, assume there is a thrackle drawing of C_4. In this drawing the edges ab and cd cross each other at a point x. Consider the closed curve c_1 obtained by starting at x, traversing ab, bc and then cd until we reach x again. Similarly, construct a closed curve c_2

by starting at x, traversing ba, ad and then dc until we reach x again. As the curves ab and cd cross at x, the closed curves c_1 and c_2 touch at x. As c_1 contains bc and c_2 contains ad we know from the definition of a thrackle that c_1 and c_2 have exactly one other point y in common at which they cross. This is a contradiction to the fact that two closed curves in the plane cross each other an even number of times (cf. 2.6).

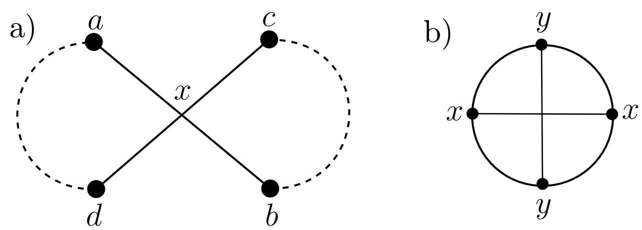

Figure 3.2. a) Apart from the crossing at x a thrackle drawing of C_4 would have to contain exactly one other crossing y between the two dashed edges, which is impossible. b) The chord diagram such a thrackle drawing would have is not cross-realizable.

Note how similar the above argument is to the one we employed in chapter 2 to prove that a double point of a crossing curve is even (see 2.7). Indeed, we can employ the results we developed there to show that C_4 cannot be thrackled. As a thrackle drawing of a circle is nothing but a drawing of a closed curve in the plane with a finite number of double-points at which the curve crosses itself, we can associate with each thrackle drawing a chord diagram. Moreover, *any* thrackle drawing of the C_4 would have the chord diagram given in Figure 3.2b): the chord diagram of a C_4-thrackle is entirely determined by the thrackle condition. It is a consequence of 2.7 that this chord diagram cannot be cross-realized as both chords are interlaced with odd many others. This observation, that thrackle drawings of circles are curves in the sense of chapter 2 and hence have a chord diagram, is the key idea of this chapter.

We have now seen that while C_3 and C_5 are thrackleable, C_4 is not. We will soon show that C_6 and, in fact, all other circles are thrackleable. But given our knowledge about C_3, C_4 and C_5 we can already observe that

3.2. Thrackleability is not closed under the (topological) minor relation.

This may appear unusual because due to Kuratowski's criterion for the planar realizability of graphs, one's first hope might have been to characterize thrackleable graphs in terms of some set of obstructions under the minor relation. Different methods will have to be employed.

3.2 Edge Duplication

We now want to decide for each C_k whether it has a thrackle drawing or not. To do this a means to construct a thrackle drawing of a larger graph from the thrackle drawing of a smaller graph would be useful. Conway introduced the concept of **edge duplication** to achieve just that. The next proof will give an example of this concept.

3.3. If C_k is thrackleable, so is C_{k+2}.

Proof. Let T be a thrackle drawing of C_k. We replace one edge a in T by a path a_1, a_2, a_3 of three edges as shown in Figure 3.3 to obtain a drawing T'. The three new edges a_1, a_2, a_3 are understood to follow the arc of the original edge a, such that the a_i cross all of the edges a crossed in T, which are all edges of T' except e and f. (Formally, this is an application of 1.9.) Additionally, a_2 and a_3 cross e while a_1 has a common endpoint with e. Similarly, a_1 and a_2 cross f while a_3 has a common endpoint with f. So, for any edge b already present in T and any new edge a_i the pair $\{a_i, b\}$ have a point in common.

On the other hand, any two edges b, c that were already present in the original drawing have a point in common in T' because they had one in T. Finally, $\{a_1, a_2\}$ and $\{a_2, a_3\}$ have a common endpoint while $\{a_1, a_3\}$ cross each other. The thrackle condition is thus satisfied for any two edges in T': we have constructed a thrackle drawing T' of C_{k+2} from a thrackle drawing T of C_k.

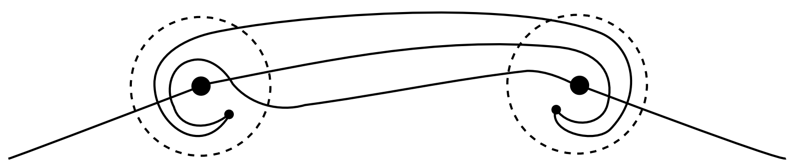

Figure 3.3. One edge is replaced with a path of three edges.

□

Before we elaborate on edge duplication, let us reap the benefits of this result.

3.4. All circles except C_4 are thrackleable.

Proof. As C_5 is thrackleable, it follows 3.3 by that all odd circles are thrackleable. C_6 can be thrackled as shown in Figure 3.4 and again we can apply 3.3 to obtain that all even circles with more than four edges are thrackleable.

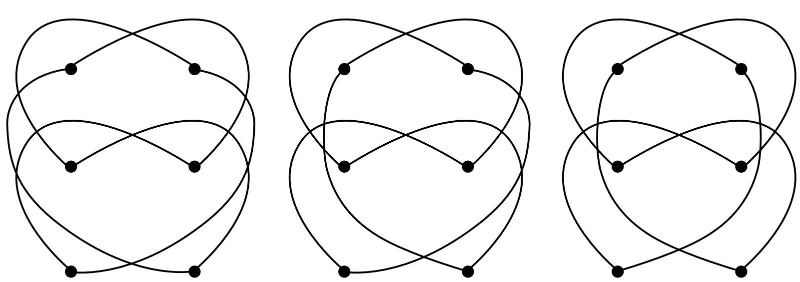

Figure 3.4. Three thrackle drawings of C_6. These drawings were obtained by edge duplicating a path of two edges (see below) and connecting the two paths in the right way. In section 3.5 we shall see that these are indeed all thrackle drawings of C_6.

□

Let us now return to edge duplication. In the literature this concept is usually explained just by looking at a concrete example such as the one given above. However, it is the experience of the author that one can easily confuse oneself when trying to check a more complex construction in such an ad-hoc manner. So he would like to invite the reader to bear with him trough the introduction of a tad more machinery which, he hopes, will greatly increase the confidence of the reader in the validity of the results presented in the next section.

The idea is that if we construct a thrackle drawing T' of G' from a thrackle drawing T of G by inserting new edges in a small neighborhood of an edge $e \in E(G)$, we know that these will cross all edges e' that cross e. As already indicated a formal proof of this statement is essentially a careful application of our "Nice Neighborhood Lemma" 1.9. A thrackle drawing of G can be interpreted a plane drawing $\tilde{d}$ of a graph $\tilde{G}$ by considering all the crossings as vertices of degree 4. Applying 1.9 to $(\tilde{G}, \tilde{d})$ we obtain nice neighborhoods of all the vertices and edges as indicated in 1.9. A neighborhood around a vertex v from our our original graph G will be called **disk** and denoted D_v. For an edge $e = vw \in E(G)$ we consider the union of all the neighborhoods of all the edges and vertices corresponding to e in $\tilde{G}$ except for D_v and D_w. We will call this union the **strip** S_e. Taking D_v, D_w and S_e together we obtain a region that in essence looks like the on depicted in Figure 3.5. Note that while all the disks are pairwise disjoint and while any strip is disjoint from all of the disks, two strips are disjoint if and only if the corresponding edges do not cross. Given this construction, we can find a curve segment in S_e with its two end points on the boundary of D_v and D_w which "runs parallel to e" in the sense that it does *not* cross e and *does* cross exactly the edges e crosses (as often as e and in exactly the same order).

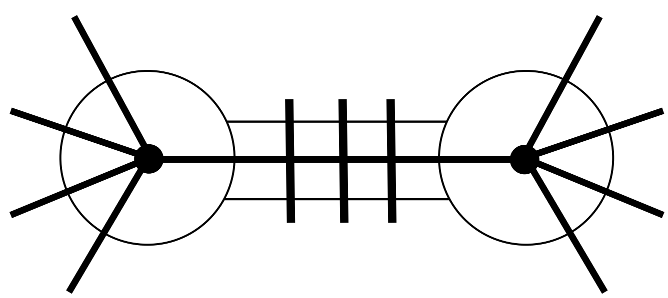

Figure 3.5. We can find a nice neighborhood around an edge e in a thrackle drawing that can be decomposed into two disks and one strip as shown above. The strip contains all the crossings of e with other edges.

This allows us to carry out explicit constructions on a drawing T without knowing what T actually looks like, as we will confine ourselves to inserting arcs that run parallel to an edge and modifying the drawing in small neighborhoods of the vertices. In particular, when we represent our constructions on a thrackle drawing T of G in the figures and sketches below, we start out with by a *different* drawing of G (a non-thrackle drawing which often is planar) and use that drawing to indicate the local modification of T.

A **local modification** of T is a drawing T' of some graph G' that differs from T only in that some disks and strips have been replaced, such that: Each edge consists of a middle part contained in a strip and two end pieces contained in the two incident disks. Each strip contains a set of middle parts that run from one disk to the other and do not have any points in common. Every edge that corresponds to one of these middle parts is said to belong to that strip. Each disk contains a set of vertices and a set of end pieces. Each end piece connects a vertex with a middle part. All end pieces are simple and if two end pieces have a common point they either end in the same vertex or they cross. And, of course, no three edges cross at a single point.

A local modification of a drawing T of G can be defined by taking any drawing of G with the same rotation system as T, indicating disks and strips around the vertices and edges and giving the local modifications of these disks and strips.

The key property of a local modification T' of T is that if e' is an edge in T' that belongs to the strip S_e of e in T and the analogous statement holds for f' and f, then

$$\#\text{crossings of the middle parts of } e' \text{ and } f' = \#\text{crossings of } e \text{ and } f$$

The following lemma tells us what we have to check to make sure that a given local modification produces a thrackle drawing T', assuming T is a thrackle drawing.

3.5. Edge Duplication Lemma

Let T' be a local modification of a thrackle drawing T. T' is a thrackle drawing if the following two conditions are satisfied:

- For every modified disk D_v and every pair of edges $\{e, f\}$ meeting D_v and belonging to *different strips*, e and f have a point of D_v in common.
- For every modified strip S and every pair of edges $\{e, f\}$ belonging to S, e and f have a common point in one of the two disks, but not in the other.

Proof. First we note that the two conditions required in the lemma hold automatically for all unmodified disks and strips. In an unmodified disk all ends are incident with a single vertex and do not cross one another. An unmodified strip contains only one middle part so there is nothing to check.

We have to check that for every two edges e', f' the thrackle condition is satisfied. As we require in the definition of a local modification that all new multi-points are of the right type (i.e. either a common endpoint at a vertex, or a crossing, etc.) and all old multi-points are of the right type because T is a thrackle, we only need to check that e' and f' have exactly one point in common. Let e and f denote the corresponding edges in T. We now consider the following cases:

e' and f' belong to the same strip S.

By construction the middle parts of e' and f' have no point in common. However, as the condition holds for S, we know that e' and f' meet in precisely one of the two incident disks.

e' and f' belong to different strips.

If e' and f' do not share a disk, we have

$$\begin{aligned} & \#\text{common points of } e' \text{ and } f' \\ = \; & \#\text{crossings of the middle parts of } e' \text{ and } f' \\ = \; & \#\text{crossings of } e \text{ and } f \\ = \; & 1 \end{aligned}$$

If e' and f' have a disk D_v in common, they do not have another disk in common, because otherwise they would belong to the same strip, as G does not have multi-edges. This implies, however, that e and f are incident and do not cross in T. So the strips of e' and f' are disjoint and thus the middle parts of e' and f' have no point in common. The common point of e' and f' are therefore the common points of e' and f' in D_v and the condition guarantees that there is precisely one of those. □

In practice this means that we can check the correctness of a local modification by looking at the disks one at a time, and then looking at the strips one at a time. The consequence is that we can effectively reuse parts of a local construction: if a modification of a disk meets the condition given above it will do so in any drawing, so we do not have to check that again. Similarly, if two neighboring disk-modifications are connected by a strip containing k edges, and we have checked that the condition is satisfied for this strip, we can connect these two disk-modifications with the given orientation by k edges in another drawing without having to check the condition for that strip again.

The duplication of a single edge is a simple application of the above lemma, which will lead us to the duplication of paths, which in turn will provide us with building blocks we will reuse extensively in the next section. Consider Figure 3.6.

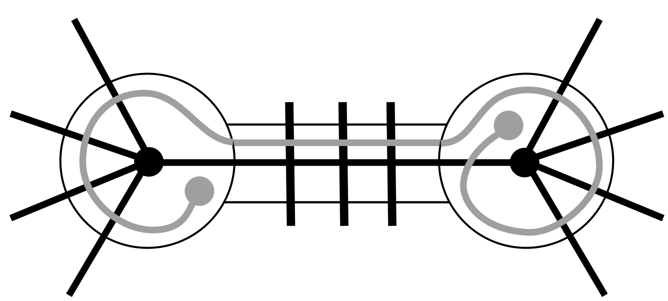

Figure 3.6. Thrackle drawing of a duplicated edge.

In the left disk, all black edges have the black vertex in common and the gray edge crosses every black edge belonging to a different strip, so the left disk satisfies the condition. The same is true for the right disk. The gray and the black edge in the strip have a common point in the right disk but not in the left. As all the other disks and strips of the drawing have not been modified, and as the construction is evidently not dependent on the number of edges incident to the black vertices, we have shown that if a graph $G = (V, E)$ has a thrackle drawing, so has $G' = (V \dot{\cup} \{v', w'\}, E \dot{\cup} \{v'w'\})$.

Note that the constructions in the left and right disks of Figure 3.6 can be combined as shown in Figure 3.7a). Again all black edges have the black vertex in common, the two gray edges have the gray vertex in common and both gray edges cross all black edges in strips different from their own. Therefore, the disk-modification shown in Figure 3.7a) satisfies the condition of the lemma. Note that the one gray edge crosses the black edge it belongs to while the other does not. This allows us to use this construction at two neighboring disks such that the condition is satisfied for the strip in between as can be seen in Figure 3.7b).

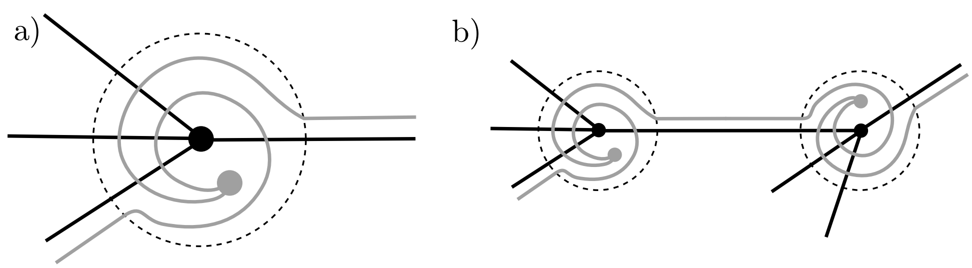

Figure 3.7. Several instances of the modified disk shown in a) can be combined as in b) such that the condition of the Edge Duplication Lemma is satisfied for both disks and the strip in between.

Note that the two gray vertices in Figure 3.7b) are placed on opposite sides of the black edge. At each disk the strips belonging to the gray edges both border on the component of $D_v \setminus T$ which the gray vertex resides in. If we repeat the construction in Figure 3.7b) along the edges the gray curve ends belong to, we proceed along a left-right-path in T. If the left-right-path is open, we can end the path (Figure 3.8) by using the appropriate disk-modification from Figure 3.6, or if the left-right-path is closed, we can close it using the disk-modification from Figure 3.7b).

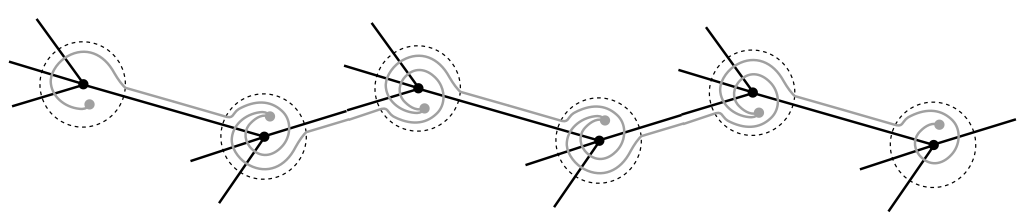

Figure 3.8. The construction in Figure 3.7 can be applied to duplicate a left-right-path.

3.6. Let T be a thrackle drawing of a graph G and $P = (\{v_1, \ldots, v_k\}, \{v_1v_2, \ldots, v_{k-1}v_k\}) \subset G$ a simple left-right path that may be closed. Then we can construct a thrackle drawing T' of $G \dot\cup P'$, $P' = (\{v_1', \ldots, v_k'\}, \{v_1'v_2', \ldots, v_{k-1}'v_k'\})$, i.e. the graph G with an additional disjoint copy of P. Note that an induced path $P \subset_{\text{ind}} G$ is always a left-right path in T.

3.3 The Thrackle Conjecture

3.7. Thrackle Conjecture CONWAY

If a graph G with n vertices and m edges is thrackleable, then $m \leqslant n$.

Or, equivalently: any graph with more edges than vertices has no thrackle drawing.

One thing that makes this conjecture interesting is the fact that it is apparently very hard to prove. Conway formulated the conjecture in the 1960s and it still remains open today (see [10], [7], [2]). Lovász, Pach and Szegedy gave an upper bound of $m \leqslant 2n - 3$ in [7] which was later improved by Cairns and Nikolayevsky to $m \leqslant \frac{3}{2}(n-1)$ in [2].

While recent research was directed at establishing upper bounds for the number of edges a thrackleable graph can have, early attempts to tackle the conjecture took a more constructive approach, making use of the concept of edge duplication. Woodall [10] showed that it suffices to establish that none of a certain class of graphs have a thrackle drawing. We will build on this result to convert the thrackle conjecture into a statement about the thrackle drawings of circles C_{2n} which are curves in the sense of chapter 2. The idea is that we can then apply our machinery for handling chord diagrams to cast the thrackle conjecture into a combinatorial form (see section 3.4) that does not take recourse to continuous concepts such as drawings.

The fact that thrackleability is closed under the subgraph relation suggests that it may be possible to prove the Thrackle Conjecture by showing that a small class of graphs cannot be thrackled: To construct a class of obstructions $\mathcal{C}$ with more edges than vertices, such that

- every graph $G = (V, E)$ with $|E| > |V|$ contains one graph in $\mathcal{C}$ as a subgraph, and
- all graphs in $\mathcal{C}$ are not thrackleable

would prove the Thrackle Conjecture.

The construction we are now going to do is slightly more involved but true to the above in spirit: Given a graph G that has more edges than vertices and assuming we have a thrackle drawing of G, we are going to construct a thrackle drawing of a graph G', which belongs to the class of 8-graphs (see below). As each 8-graph has more edges than vertices, the Thrackle Conjecture is thus equivalent to the conjecture that no 8-graph can be thrackled.

An **8-graph** is a graph that consists of two circles C_i and C_j that have precisely one vertex in common. We are going to denote these graphs with $8_{i,j}$. Note that $8_{i,j}$ has $i + j - 1$ vertices and $i + j$ edges. Hence, if any $8_{i,j}$ has a thrackle drawing, the Thrackle Conjecture is false. Following Woodall [10], we are going to show that the converse is also true.

3.8. Thrackle Conjecture – 8-Graph Version

None of the graphs $8_{2k,2k}$ for $k \geqslant 3$ has a thrackle drawing.

Before we can prove that the 8-Graph Version is indeed equivalent to the Thrackle Conjecture, we need to learn more about thrackle drawings of 8-graphs, should they exist. The definition of a thrackle prescribes how often the two circles of an 8-graph cross each other in a thrackle drawing. The computation of this number immediately yields the following result (see [7]).

3.9. In a thrackle drawing of $8_{i,j}$, the two circles cross each other at their common vertex if and only if both i and j are odd.

Proof. In the drawing of $8_{i,j}$ both circles C_i and C_j are closed curves and thus have to cross each other an even number of times. Hence, they will cross at the common vertex v iff their edges cross an odd number of times. Every edge C_i crosses every edge of C_j with the exception of the edges incident to v: the two edges of C_i incident to v do not cross the two edges of C_j incident to v. The number of crossings of edges of C_i with edges of C_j therefore is $ij - 4$, which is odd iff both i and j are odd. □

We can now show the equivalence of the two versions of the thrackle conjecture.

Proof of the equivalence of the Thrackle Conjecture and its 8-Graph Version. All that is left to show is that given a thrackle drawing T of a graph $G = (V, E)$ with $|E| > |V|$, we can construct a thrackle drawing of some $8_{2k,2k}$. As G has more edges than vertices there has to be a connectivity component of G with the same property, so without loss of generality, we can assume that G is connected. Let n be the number of vertices of G and S a spanning tree. S has $n - 1$ edges and we know $|E| \geqslant n + 1$ so we can pick two edges $e, e' \in E$ that are not in S. In S, there is a unique path P from one vertex of e to the other, giving us a circle $C = P \cup e$. Similarly we obtain a circle $C' = P' \cup e'$.

Now, the intersection of P and P' must again be a (possibly empty) path, it cannot be disconnected: If v and w are vertices in the intersection, and $U \subset P$ and $U' \subset P'$ simple paths connecting the two, then $U = U' \subset P \cap P'$, for otherwise $U \cup U' \subset T$ would contain a circle. Given that $e \neq e'$, we can conclude that the relationship between C and C' must be in one of the following (see Figure 3.9):

1. C and C' have neither an edge nor a vertex in common, but there is a path P connecting the two, i.e. $C \cup P \cup C'$ is a "dumbbell graph".

2. C and C' meet at exactly one vertex, i.e. $C \cup C' = 8_{i,j}$ for some i, j.

3. The intersection of C and C' is a path, i.e. $C \cup C'$ is a Θ-graph.

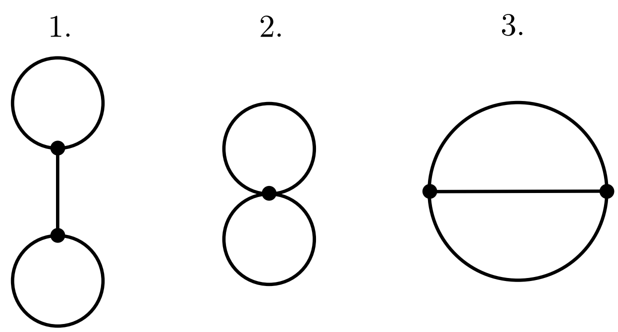

Figure 3.9. A counterexample to the thrackle conjecture always contains a subgraph that is a subdivision of one of the three (multi)graphs shown above.

In each case we have to construct a thrackle drawing of a $8_{2k,2k}$ from the given thrackle drawing of $C \cup C'$ (resp. $C \cup P \cup C'$). Note that using the construction employed in the proof of 3.3 we can construct a $8_{2k,2k}$ -thrackle from a $8_{2i,2j}$-thrackle. A $8_{2i,2j}$-thrackle in turn can be constructed in each of the three cases above using the following constructions:

i. constructing a $8_{i,j}$-thrackle from a Θ-thrackle

ii. constructing a $8_{i,j}$-thrackle from a dumbbell-thrackle

iii. constructing a $8_{i,2j}$-thrackle from a $8_{i,j}$-thrackle, if i is even and j is odd

iv. constructing a $8_{i,j+i}$-thrackle from a $8_{i,j}$-thrackle, if both i and j are odd

Note that we have to treat iii. and iv. separately, because we know from 3.9 that if both circles are odd we have a different local situation at the common vertex. All that is left to show, is that these constructions are indeed possible.

Constructions i. and ii.

We will treat the cases of a dumbbell and a Θ-graph together by giving a local modification of the path: in both cases we have a subgraph of the form shown in Figure 3.10a). Note that because there are no loops or multi-edges in G we can assume without loss of generality that the edges a, b, c, d in Figure 3.10b) are distinct. We then use the local modification given in Figure 3.10b).

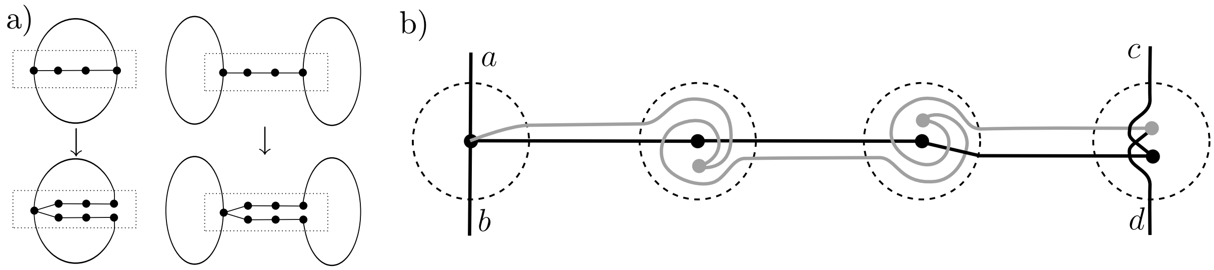

Figure 3.10. a) shows how we can apply the local modification given in b) to deal with cases i. and ii.

The modified disk on the left satisfies the condition because all ends are incident to one vertex. The disk on the right satisfies the condition of the Edge Duplication Lemma, because c has exactly one point in common with each of the other edges, and by symmetry so has d. Note that the two edges that belong to one strip have no point in common. The disks in between satisfy the condition as we know from the previous section. In each disk D_w the pair of edges coming from the left have no common point in D_w, while the pair of edges leaving on the right do have a common point. Thus, the condition is fulfilled for every strip, even if the path consists of only a single edge. The graph given by the drawing we just constructed is in both cases an $8_{i,j}$.

Construction iii.

Let C_i and C_j denote (the thrackle drawings of) the two circles of the $8_{i,j}$-thrackle. Let v_0 be their common vertex and $C_j = v_0v_1...v_{j-1}$. The local modification is given in Figure 3.11. We start out by edge-duplicating the open path $v_0v_1...v_{j-1}$ such that the new vertex in D_v is on the right-hand side of $S_{v_0v_1}$ and the two edges belonging to $S_{v_0v_1}$ do not cross. As the path has $j-1 \equiv 0 \pmod 2$ edges and as the side on which the new vertex is placed alternates with each edge we duplicate, the new vertex in $D_{v_{j-1}}$ is placed on the right-hand side of $S_{v_{j-1}v_0}$. Another consequence of this construction is that the two edges belonging to $S_{v_{j-2}v_{j-1}}$ cross at $D_{v_{j-1}}$. The situation after this first step of the construction is shown in Figure 3.11a).

We now modify $D_{v_{j-1}}$, $S_{v_{j-1}v_0}$ and D_{v_0} as shown in Figure 3.11b). These three satisfy the condition of the Edge Duplication Lemma by construction. We note that, as in Figure 3.11a), the edges belonging to $S_{v_0v_1}$ do not cross in D_{v_0} and the edges belonging to $S_{v_{j-2}v_{j-1}}$ do cross in $D_{v_{j-1}}$, so $S_{v_0v_1}$ and $S_{v_{j-2}v_{j-1}}$. All other disks satisfy the condition, because they were not changed, and all other strips satisfy the condition because the incident disks were not changed. So, by the Edge Duplication Lemma, Figure 3.11b) does really define a thrackle drawing. The path of gray edges and the path of black edges, which are both of the same length, meet at v_0 and at v_0', so they form an even circle and we have indeed constructed a thrackle drawing of $8_{i,2j}$.

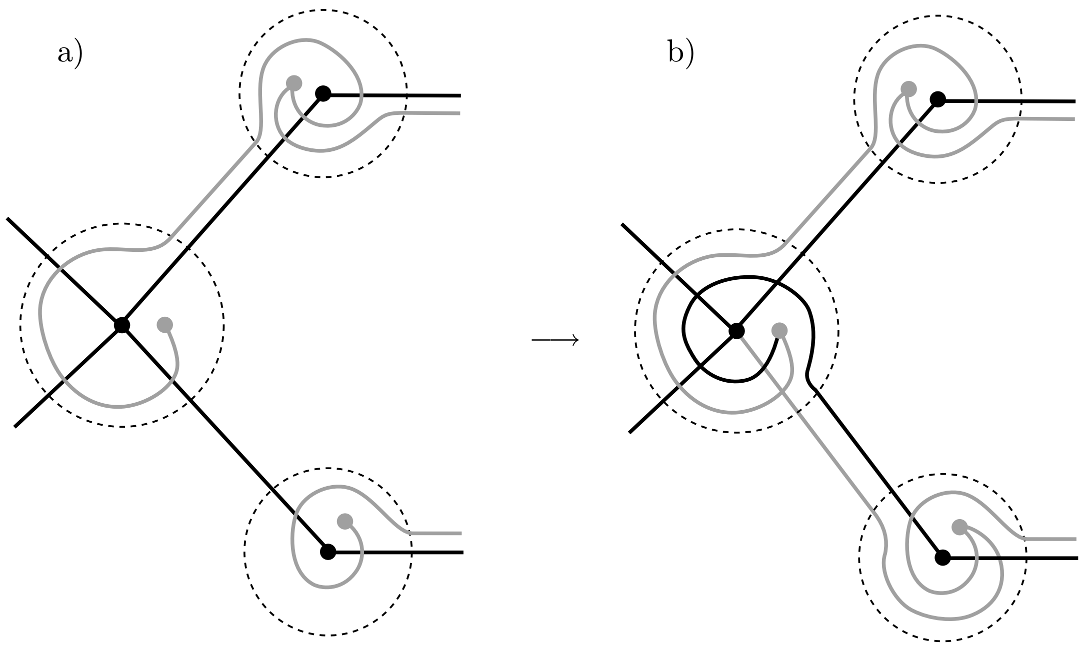

Figure 3.11. a) The first step is a path duplication. Because there are an even number of disks between v_1 and v_{j-1}, the gray edge leaves D_{v_1} and arrives at $D_{v_{j-1}}$ to the right of the black edge. b) This further modification yields a $8_{i,2j}$ thrackle. Note that both edges belonging to $S_{v_{j-1}v_0}$ are modified.

Construction iv.

Let C_i and C_j denote the circles of $8_{i,j}$ and v_0 their common vertex. Both i and j are odd, so C_i and C_j cross at v_0. This means we cannot use construction iii. We are still going to duplicate the edges of one circle, say $C_i = v_0 ... v_{i-1}$, but instead of forming a larger circle from the new edges and the edges of C_i we will form a larger circle from the duplicate of C_i and C_j to obtain a thrackle drawing of $8_{i,j+i}$. As i and j are both odd C_{j+i} is even as required.

See Figure 3.12 for the construction we will use. C_i is drawn vertically (and in gray) and C_j is drawn horizontally. Again, it is easiest to think of this as a two step construction. First, we duplicate the path $v_1 ... v_{i-1}$ such that the new vertex v_1' is on the left-hand side of $S_{v_1v_2}$ and such that the two edges belonging to $S_{v_1v_2}$ cross each other in D_{v_1}. The duplicate is shown in black. Accordingly, as there are an odd number of strips $S_{v_1v_2}, ..., S_{v_{i-2}v_{i-1}}$ between D_{v_1} and $D_{v_{i-1}}$, the new vertex v_{i-1}' in $D_{v_{i-1}}$ is on the right-hand side of $S_{v_{i-1}v_0}$ and the edges belonging to $S_{v_{i-2}v_{i-1}}$ cross each other in $D_{v_{i-1}}$ (see Figure 3.12a). The second step is to modify D_{v_0}, D_{v_1} and $D_{v_{i-1}}$ as shown in Figure 3.12b). Each of these three disks satisfy the condition of the Edge Duplication Lemma as do the strips incident with them, so Figure 3.12b) does define a thrackle drawing, and it is indeed a thrackle drawing of $8_{i,j+1}$: starting from v_0 to the right, we arrive at D_{v_0} from the left. Given the local modification of D_{v_0} we continue downwards along the gray edge only to arrive at v_0 again from above. This gives us C_{j+i}, the vertical black edges give us C_i and both circles meet at v_0 where they do not cross each other.

The disks of C_j are not modified, except for D_{v_0}. The local modification of D_{v_0} is new, but evidently the condition is fulfilled. The local modifications of the remaining disks are known, so we do not have to check them. To apply the lemma, we still need to check the strips of C_i.

The two edges belonging to S_1 have a common vertex in D_{v_0} but do not meet in the other disk. For S_2, ..., S_{i-1} we employed the usual path duplication so we do not need to check them. It is correct that the gray edge is on the same side of the black edge in S_1 and S_i as there are even many of the disks known from path duplication in between these two strips. Also, we know that the black and gray edges belonging to S_i do not cross in $D_{v_{i-1}}$. Applying the Edge Duplication Lemma we conclude that Figure 3.12 gives us a thrackle drawing.

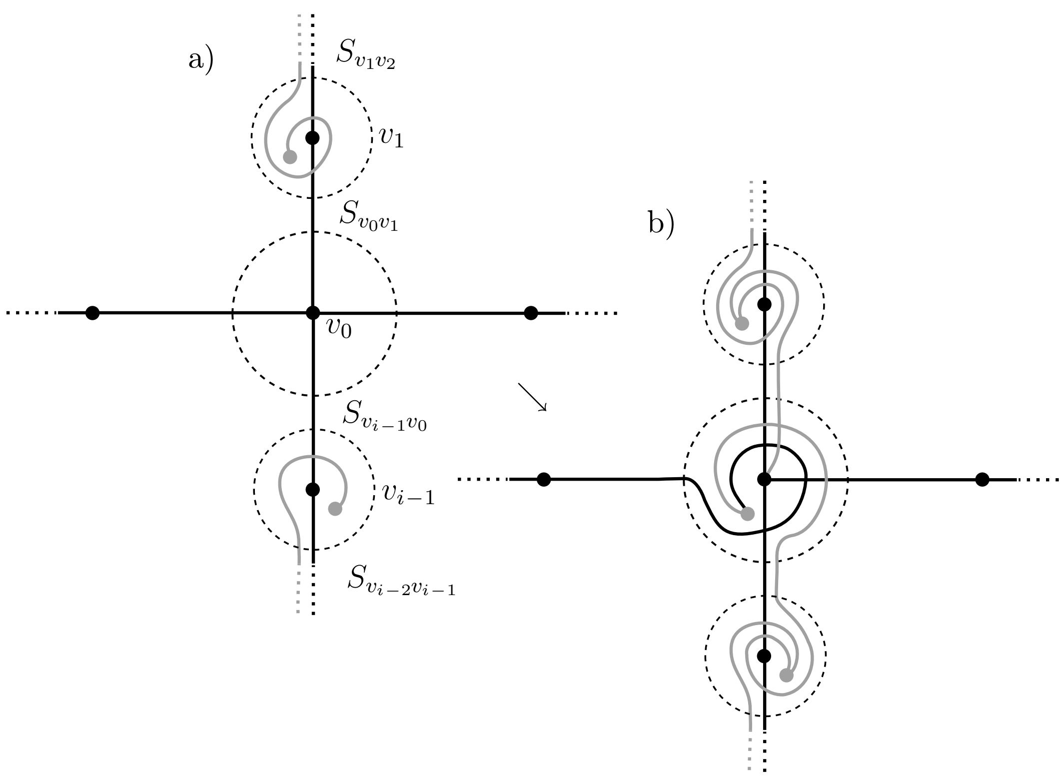

Figure 3.12. a) The first step is to duplicate one of the two circles (minus the two edges incident with v_0). As there are an even number of disks between the two modified disks, the gray edge starts out at the top as it arrives at the bottom: to the left of the black edge. b) In the second step we modify the three disks as shown to obtain a thrackle drawing of $8_{i,j+i}$.

□

3.4 From Thrackles to Curves

As we know the thrackle conjecture is true if and only if all $8_{2i,2i}$ are thrackleable. A thrackle drawing of a $8_{2i,2i}$-graph is a closed curve in the plane, so we can apply the machinery we have developed for the study of closed curves in the plane. As both circles of $8_{2i,2i}$ are even, the common vertex v is a touching point of the two. Consequently, the two Euler tours of $8_{2i,2i}$ give us a closed curve that touches itself at v and a closed curve that crosses itself at v, respectively.

Moreover, we can also construct a thrackle drawing of a circle C_{4i} from the thrackle drawing of $8_{2i,2i}$: We modify a small neighborhood around v as shown in Figure 3.13. Any two edges have a common point in this disk, so the drawing we obtain is really a thrackle drawing of a circle.

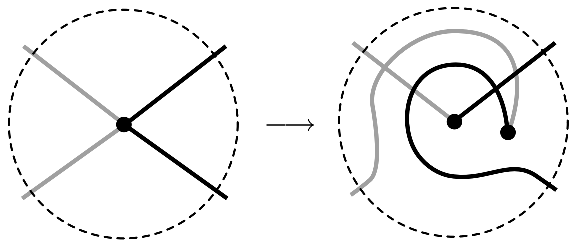

Figure 3.13. Local modification to apply to a neighborhood of the common vertex v of a 8_{2i2j}-thrackle to obtain a C_{2i+2j}-thrackle. The half-edges belonging to the one circle are shown in gray while those belonging to the other are shown in black.

We have now seen that we can interpret a $8_{2i,2i}$-thrackle as a crossing curve or that we can apply a local modification to a $8_{2i,2i}$-thrackle to obtain a particular C_{4i}-thrackle which again is a crossing-curve. How do the chord diagrams and interlacement graphs of these two curves look? To answer this question we will first consider chord diagrams of circle thrackles in general.

Chord Diagrams of Circle Thrackles

Let T denote a thrackle drawing of a circle C_i. As it is a drawing of a circle, T is a closed curve with a chord diagram C. An edge of C_i corresponds to an interval on the rim of C. Identifying the rim with S^1 and interpreting it as the preimage of the curve T, the intersection of two intervals is a single point if the corresponding edges are incident and empty otherwise. Interpreting the intervals as vertex sets of C, the intervals are pairwise disjoint and cover $V(C)$. The chords in C are exactly the crossings of two edges of C_i. As T is a thrackle drawing, we conclude that if two edges are incident, there is no chord between their intervals, and if they are incident, there is exactly one chord between their intervals. We do not know the order of the end-points of the chords in a given interval.

We already know which circles have thrackle drawings, but we do not know in which order a given edge will cross the other edges in such a drawing. In this regard it is helpful to formulate a characterization of circle thrackles in terms of chord diagrams.

3.10. A circle has a thrackle drawing, if there exists some permutation of the chord-ends in the intervals of the chord diagram as given above, such that the resulting chord diagram is realizable. See Figure 3.14a).

The kind of problem we are trying to solve is this: Given a chord diagram along with a partition of the chord-ends into disjoint intervals, is there some permutation of the chord-ends in each interval such that the resulting chord diagram is realizable? Obviously the answer to the above question does not depend on the order of the chord-ends within an interval of the chord diagram we are given. Thus, an equivalent and more compact way of asking the same question is: Is there some cross-realizable chord diagram that can be transformed into 3.14b) by contracting edges on the rim? Note that in this context we allow the process of edge contraction to create loops and multi-edges (they are not removed after a contraction).

More generally, the contraction of a chord diagram of a C_n-thrackle is a K_n with a $C_n \subset K_n$ as rim.

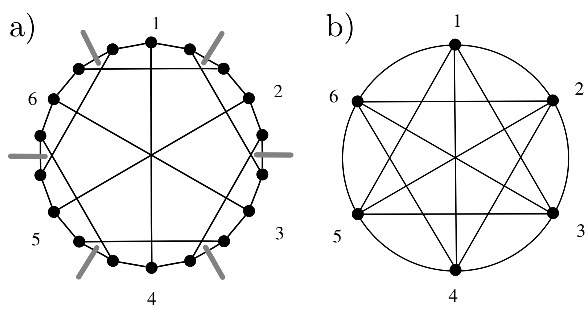

Figure 3.14. a) A thrackle drawing of the C_6 gives a permutation of the curve ends in each interval of the above chord diagram, such that the resulting chord diagram is cross-realizable. b) Given a chord diagram of a thrackle drawing of C_6, contract the edges on the rim that fall within an interval and you obtain the K_6.

In Figure 3.15a) and b) the chord diagrams of the three C_6-thrackles we already met are given along with their chord diagrams. In the next section we are going to use an equational version of the Rosenstiehl Criterion 2.15 for crossing curves to show that these are indeed all C_6-thrackles (up to equivalence of curves).

Interlacement Graphs of Circle Thrackles

An interval I_i on the rim of a chord diagram determines a set of vertices S_i in the interlacement graph: a chord v is in S_i iff one of its end-points is in I_i. A partition of the chord-ends into disjoint intervals gives a covering of the vertex set of the interlacement graph, in which every vertex is in at least one and at most two sets. In the case of circle thrackles the contracted chord diagram has now loops and multi-edges, which means in particular that every vertex in the interlacement graph is in two sets S_i and S_j with $i \neq j$. If we do not know the order of the chord-ends in an interval I_i, we do not know for two vertices v, $w \in S_i$ whether they are adjacent in the interlacement graph. However, the induced subgraph on the vertex set S_i has to be a "permutation graph" (see below). For two vertices v, w that are not in any common set, the adjacency of v and w is determined by the contracted chord diagram.

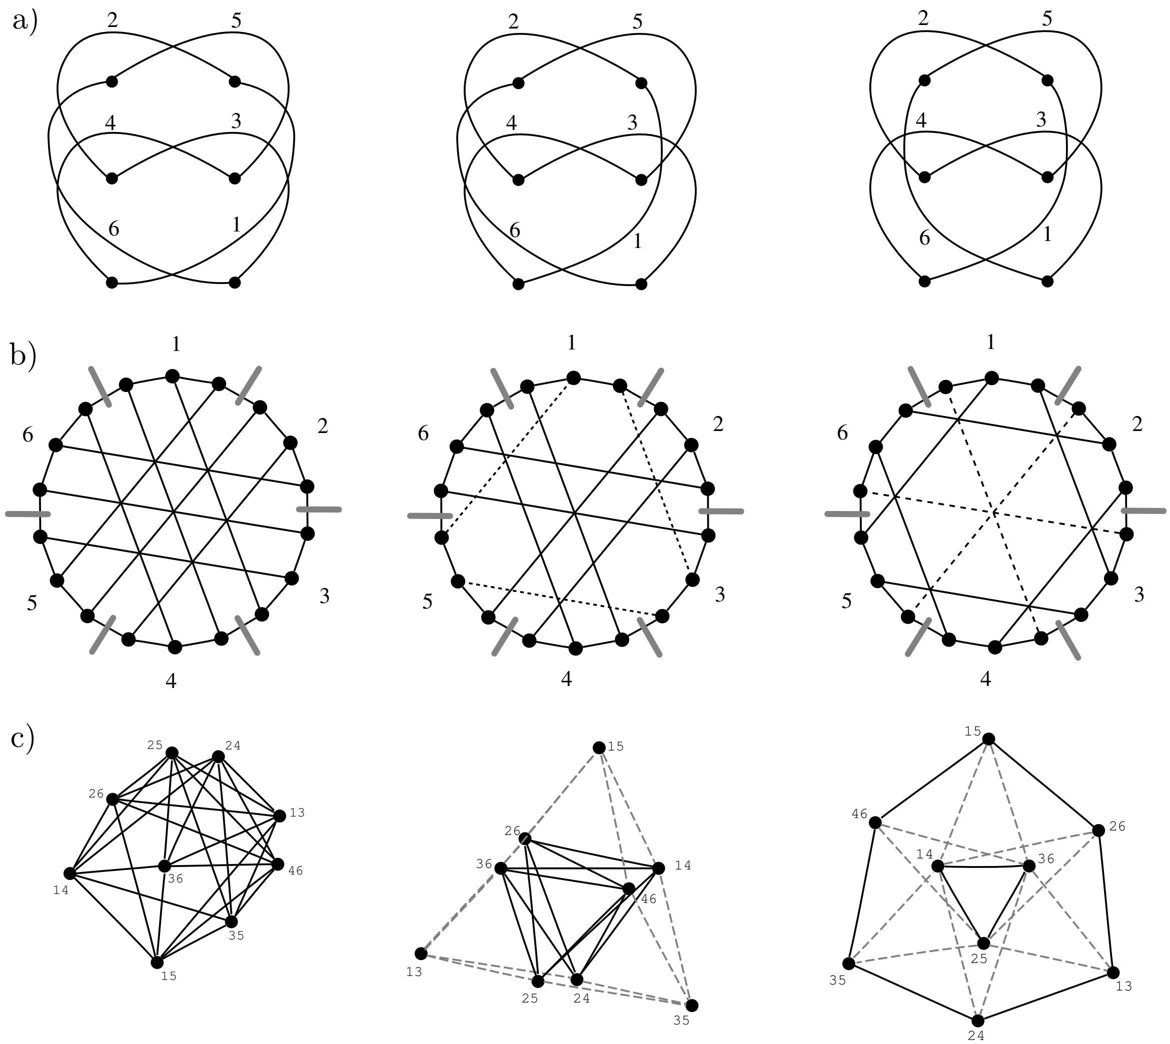

Figure 3.15. We have a) the three C_6-thrackles in the first row, b) their corresponding chord diagrams in the second row and c) their interlacement graphs in the last row. The classes of the chords are given by the solid resp. the dashed lines in the chord diagram. In the interlacement graphs, the dashed lines mark even edges, i.e. edges between vertices that have an even common neighborhood.

A **permutation graph** is a graph $G=(V,E)$ for which there exists a linear order $\prec$ on V and a permutation π of V such that $v,\, w \in V$ are adjacent in G iff they are inverted with respect to $\prec$ and π. Suppose that V is a subset of the vertices of an interlacement graph with an associated chord diagram C. If V corresponds to an interval on the rim of C, the linear order $\prec$ is given by C. Note that while the linear order is determined only up to reversal, the permutation graph is of course uniquely determined as the permutation graph of $(V,\prec,\pi)$ is identical to the permutation graph of $(V,\prec^{-1},\pi)$.

In general the structure of the interlacement graph $\Lambda = (V,\, E)$ of a C_k-thrackle can be described as follows.

$$V = \{\{i,j\}\colon 1 \leqslant i < j \leqslant k \text{ and } i, j \text{ are not next to each other}\}$$

We abbreviate the unordered pair $\{i, j\}$ as ij. The requirement that i, j are not "next to each other" can be formalized as $i+1<j$ and $i=1 \Rightarrow j \leqslant k-1$. Now, whether two vertices ij and ab are adjacent or not is uniquely determined by the thrackle condition iff all four integers i, j, a, b are distinct. In that case, the vertices ij and ab are adjacent if the pairs of numbers are interlaced. All other adjacencies are undetermined. However, we can observe that the induced subgraph on each of the sets $P_i = \{ij \in V\}$ is a permutation graph. On each of the P_i, the order is given by the rule $ij \prec i(j+1)$ where addition is performed in $\mathbb{Z}_k+1$. This structure can be visualized as shown in Figure 3.16.

													14,16
												13,15	13,16
											12,14	12,15	12,16
										11,13	11,14	11,15	11,16
									10,12	10,13	10,14	10,15	10,16
								9,11	9,12	9,13	9,14	9,15	9,16
							8,10	8,11	8,12	8,13	8,14	8,15	8,16
						7,9	7,10	7,11	7,12	7,13	7,14	7,15	7,16
					6,8	6,9	6,10	6,11	6,12	6,13	6,14	6,15	6,16
				5,7	5,8	5,9	5,10	5,11	5,12	5,13	5,14	5,15	5,16
			4,6	4,7	4,8	4,9	4,10	4,11	4,12	4,13	4,14	4,15	4,16
		3,5	3,6	3,7	3,8	3,9	3,10	3,11	3,12	3,13	3,14	3,15	3,16
	2,4	2,5	2,6	2,7	2,8	2,9	2,10	2,11	2,12	2,13	2,14	2,15	2,16
1,3	1,4	1,5	1,6	1,7	1,8	1,9	1,10	1,11	1,12	1,13	1,14	1,15	

Figure 3.16. This figure shows all the vertices of an interlacement graph of a C_{16} thrackle. The chord $6, 12$ that corresponds to the chord that connects intervals I_6 and I_{12} is singled out. All vertices with a white background are not adjacent to 6, 12 in any interlacement graph of a C_6 thrackle. The vertices with a dark gray background are adjacent to 6, 12 in every interlacement graph of a C_6 thrackle. The vertices with a light gray background may or may not be interlaced with $6, 12$, depending on the order of the chords in the intervals I_6 and I_{12} respectively.

The Thrackle Conjecture in Terms of Chord Diagrams

If a graph G on k vertices with a given rim $C_k \subset G$ can be obtained from a chord diagram C by contracting edges on the rim (without deleting multi-edges or loops), we write $G < C$ and say G is a contract of C.

Now we can provide two combinatorial reformulations of the thrackle conjecture, by simply transferring the observations made at the beginning of this section into the language of chord diagrams. That these two reformulations suffice to show the 8-graph version is immediate and that they are necessary is not much more difficult show. In the interest of presenting formally sound arguments, proofs will nonetheless be given further below.

3.11. Thrackle Conjecture – Chord Diagram Version

For all $n \geqslant 3$ there is no cross-realizable chord diagram C with $K'_{4n} < C$.

Here K'_{4n} is obtained from K_{4n} by adding a chord and removing four others as shown in Figure 3.17.

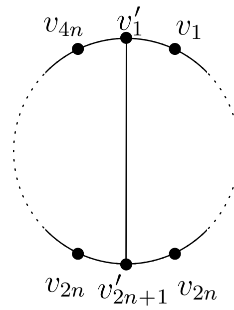

Figure 3.17. A small section of K'_{4n}. v'_1 and v'_{2n+1} are adjacent to none of the vertices that are not shown, while $\{v_{4n}, v_1, v_{2n}, v_{2n+1}\}$ are adjacent to all of the other vertices. Formally K'_{4n} can be defined as follows: Let $K_{4n} = (V, \binom{V}{2})$ with $V = \{v_1, \ldots, v_{4n}\}$ such that the rim of K_{4n} is $v_1v_2, \ldots, v_{4n-1}v_{4n}, v_{4n}v_1$. Then we define

$$K'_{4n} := (V \cup \{v'_1, v'_{2n+1}\}, \binom{V}{2}) \setminus \{v_{4n}v_{2n}, v_{4n}v_{2n+1}, v_1v_{2n}, v_1v_{2n+1}\} \cup \{v'_1v'_{2n+1}\}$$

The equivalence of the 8-graph version and the chord diagram version of the Thrackle Conjecture is immediate.

3.12. Thrackle Conjecture – Circle Thrackle Version

For all $n \geqslant 3$ there is no cross-realizable chord diagram C with $K_{4n} <$ C such that the intervals I_{4n}, I_1, I_{2n}, I_{2n+1} of C are connected as shown in Figure 3.18.

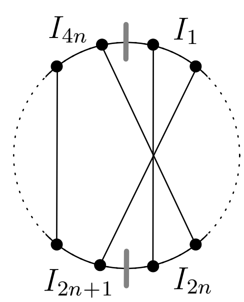

Figure 3.18. A small section of C. The structure of the chord diagram induced by the construction in Figure 3.13. All other chords have their end-points on the dashed segments of the chord diagram. A consequence is that if any other chord is interlaced with one of these four, it is interlaced with all of them: the four chords have the same neighborhood.

Proof of the Equivalence of the 8-Graph and the Circle Thrackle Version.
That the circle thrackle version suffices to prove the thrackle conjecture is immediate. We will now show that it is also necessary: Let c be a cross-realization of the chord diagram C

described in 3.12. Let c_1 and c_2 be those segments of c that correspond to the segments of the rim of C that are solid in Figure 3.18. Consider the union of the arcs $c_1 \cup c_2$ as a multigraph drawing on the sphere and put $r := c \setminus c_1 \cup c_2$. All of r is contained in one face F of $c_1 \cup c_2$ as r does not cross $c_1 \cup c_2$ and hence $c_1 \cup c_2$ is contained in $S^2 \setminus F$ which is contractible. Contracting $S^2 \setminus F$ then gives us a thrackle drawing of $8_{2n,2n}$ on the sphere, which in turn defines a thrackle drawing of $8_{2n,2n}$ in the plane. □

Note that these versions of the thrackle conjecture are statements about discrete objects. While both the original thrackle conjecture and the 8-graph version state that there do not exists certain graph drawings, which are essentially continuous functions, these two versions state that certain chord diagrams do not exist. There are only countably many objects to check and we have met many criteria to check a given chord diagram for cross-realizability. Of course that does not necessarily mean that any of these two combinatorial versions is easier to decide than the original thrackle conjecture.

The one property that makes the two formulations above hard to decide, is the fact that whether or not a given chord diagram C with $K_{4n} < C$ is cross-realizable depends heavily on the order of chords in a given interval. There are a huge number of chord diagrams C with $K_{4n} < C$ but only a few of them are realizable. This puts the difficulty of proving the thrackle conjecture into perspective: whether a given graph has a thrackle drawing or not depends heavily on the *order* in which a given edge crosses the other edges it is not incident with.

However, precisely the fact that there are few cross-realizable chord diagrams C with $K_{4n} < C$ gives hope that it might be possible to characterize all chord diagrams of circle thrackles. Given such a characterization one could perhaps conclude that the substructure given in Figure 3.18 cannot occur in the chord diagram of a circle thrackle. We will take a first step in this direction by giving a proof that the chord diagrams given in Figure 3.15 are indeed all chord diagrams of a C_6-thrackle in the next section.

A few words on computational approaches to the problem: it is possible to enumerate all (equivalence classes under isomorphy of) chord diagrams $C > K_6$ and check them for realizability on a desktop computer using a reasonably efficient implementation. To do the same for K_8 would take many years, and the first interesting case of K_{12} is utterly hopeless. However, it may be possible to obtain complete lists of chord diagrams of C_k-thrackles for larger k using smarter algorithms such as the one we will use in the next section to calculate the list in the case of $k = 6$ by hand. Also, it might be worth a try to search for a counterexample to the thrackle conjecture.

3.5 C_6-Thrackles

3.13. Figure 3.15 shows all C_6-thrackles.

Proof. We first note that any chord diagram C with $C_{K_6} < C$ has a connected interlacement graph and by 2.33 at most one cross-realization. So all we need to do is to show that the three chord diagrams given in Figure 3.15 are the only cross-realizable chord diagrams C with $C_{K_6} < C$. To this end, we are going to use the equational version of our criterion for the realizability of augmented chord diagrams 2.23. In fact, as we are dealing with crossing curves, an equational version of the original Rosenstiehl Criterion suffices, which states that the congruence

$$\langle N_x, N_y \rangle + i_{x,y} \cdot [x \text{ and } y \text{ are of the same class}] \equiv 0$$

holds for every pair of chords x, y, which reduces to

$$\langle N_x, N_x \rangle \equiv 0$$

in the case of $x = y$. Recall that in the context of 2.23 we defined a boolean predicate

$$i_{a,b} = [\text{chords } a \text{ and } b \text{ are interlaced}]$$

In the context of circle thrackles, the value of $i_{a,b}$ is fixed for chords a and b that do not share an interval and it is variable if a and b do share an interval. We are thus going to group the variables $i_{a,b}$ by the interval the two chords share. Note that $i_{a,b}$ and $i_{b,a}$ are, of course, one and the same. For each interval I_k we define a vector

$$v^k = \begin{pmatrix} i(k(k+3), k(k+4)) \\ i(k(k+2), k(k+4)) \\ i(k(k+2), k(k+3)) \end{pmatrix}$$

where the addition is in $\mathbb{Z}_6 + 1$. In our case of C_6-thrackles this gives, e.g.

$$v^1 = \begin{pmatrix} i_{14,15} \\ i_{13,15} \\ i_{13,14} \end{pmatrix}, \; v^3 = \begin{pmatrix} i_{36,31} \\ i_{35,31} \\ i_{35,36} \end{pmatrix}, \; v^5 = \begin{pmatrix} i_{52,53} \\ i_{51,53} \\ i_{51,52} \end{pmatrix}$$

Given v^k, the permutations of the chords ending in I_k is uniquely determined. Note that because $i_{kj,k\ell} = 1$ iff the chords kj and kl are inverted in the permutation, the cases

$$v^k = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} \text{ and } v^k = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$$

cannot occur: in the former case both of the "outer" chords would be inverted with the one in the middle, but not with each other, which is impossible, and in the latter case both of the "outer" chords would be inverted with each other, but not with the one in the middle, which, again, is impossible.

There is no cross-realizable chord diagram C with $K_6 < C$ and $v^1 = (0,0,0)$ or $v^1 = (1,1,1)$.

Suppose C were such a chord diagram. Then we have

$$\begin{aligned} 0 &\equiv \langle N_{13}, N_{13}\rangle \\ &\equiv 1 + \underbrace{i_{13,14} + i_{13,15}}_{\equiv 0 \text{ by assumption}} + i_{31,36} + i_{31,35} \end{aligned}$$

which means that $i_{31,35}$ and $i_{31,36}$ have to differ. This means that the first and middle values of v^3 differ, i.e. $v_1^3 \not\equiv v_2^3$. By the above argument, we conclude that $v_2^3 \equiv v_3^3$ i.e. $i_{35,36} \equiv i_{31,35}$. Considering 15 we obtain $i_{51,53} \equiv i_{52,53}$. Now we have

$$\langle N_{35}, N_{35}\rangle \equiv 1 + \underbrace{i_{35,36} + i_{35,31}}_{\equiv 0} + \underbrace{i_{53,52} + i_{53,51}}_{\equiv 0} \equiv 1$$

which means that the condition of the Rosenstiehl Criterion does not hold for the chord 35.

The key idea here is that we have

$$\begin{aligned} \langle N_{13}, N_{13}\rangle &\equiv 1 \quad + v_2^1 + v_3^1 + v_1^3 + v_2^3 \\ \langle N_{35}, N_{35}\rangle &\equiv 1 \quad + v_2^3 + v_3^3 + v_1^5 + v_2^5 \\ \langle N_{15}, N_{15}\rangle &\equiv 1 + v_1^1 + v_2^1 \quad + v_2^5 + v_3^5 \end{aligned}$$

which we will, again, exploit in the next step.

If C is a realizable chord diagram with $K_6 < C$ and $v^1 = (1,0,0)$ or $v^1 = (0,1,1)$, then $v^1 = v^3 = v^5$.

Applying 2.23 to each of the chords 15 and 13 as above, we find that $v_2^5 \equiv v_3^5$ while $v_1^3 \not\equiv v_2^3$. Applying 2.23 to the pair $15, 13$ we obtain

$$\begin{aligned} i_{15,13} \cdot [x \text{ and } y \text{ are of the same class}] &\equiv \langle N_{15}, N_{13}\rangle \\ &\equiv 1 + i_{31,36} + i_{51,52} + i_{14,15} \cdot i_{13,14} + i_{51,53} \cdot i_{31,35} \\ &\equiv 1 + v_1^3 + v_3^5 + \underbrace{v_1^1 \cdot v_3^1}_{\equiv 0} + v_2^5 \cdot v_2^3 \\ &\equiv v_2^3 + v_2^5 + v_2^5 \cdot v_2^3 \end{aligned}$$

If $v^1 = (1,0,0)$ then 13 and 15 are not interlaced and the above equation holds iff $v_2^3 \equiv v_2^5 \equiv 0$ in which case $v_3^5 \equiv 0$ and $v_1^3 \equiv 1$. From our first observation we obtain $v_2^5 \equiv 0$ and $v_2^3 \equiv 0$. Because $v^3 = (1, 0, 1)$ is forbidden, we have $v_3^3 \equiv 0$ and $v_1^5 \equiv 1$ follows from $\langle N_{35}, N_{35}\rangle \equiv 0$. So we conclude $v^1 = v^3 = v^5$.

If $v^1 = (0,1,1)$ and 13 and 15 are not of the same class, we find $v^3 = v^5 = (1,0,0)$ as above, which is a contradiction because applying the above argumentation with v^3 in the role of v^1 yields $v^1 = (1,0,0)$.

If $v^1 = (0,1,1)$ and 13 and 15 are of the same class, the equation reduces to $1 \equiv v_2^3 + v_2^5 + v_2^5 \cdot v_2^3$, which holds iff not both v_2^3 and v_2^5 are zero. Suppose $v_2^3 \not\equiv v_2^5$, then we infer $v_3^3 \equiv v_1^5$ from $\langle N_{35}, N_{35}\rangle \equiv 0$. On the other hand we have $v_1^3 \equiv v_3^5$. If $v_1^3 \equiv v_3^3$, either v^3 or v^5 would be one of the forbidden vectors $(1, 0, 1)$ or $(0, 1, 0)$. If $v_1^3 \not\equiv v_3^3$, either v^3 or v^5 would equal

$(1,0,0)$ as neither can be $(0,0,1)$ and we would again obtain a contradiction. So $v_2^3 \equiv v_2^5 \equiv 1$ and from $\langle N_{35}, N_{35}\rangle \equiv 0$ we infer that $v^1 = v^3 = v^5$.

Conclusion

By symmetry we obtain the analogous result for $v^1 = (0,0,1)$ or $v^1 = (1,1,0)$. We can now state that in a realizable chord diagram C with $K_6 < C$ we have $v^1 = v^3 = v^5 =: v$ and $v^2 = v^4 = v^6 =: v'$ where $v, v' \in \{(0,0,1),(1,1,0),(1,0,0),(0,1,1)\}$. The chord diagrams in Figure 3.15 are, respectively, of the types

- $v = v' = (1,1,0)$
- $v = (1,0,0)$ but $v' = (1,1,0)$
- $v = v' = (1,0,0)$

Considering symmetries, all that is left to show is that the three chord diagrams with

1. $v = (0,0,1)$ and $v' = (1,0,0)$
2. $v = (0,1,1)$ and $v' = (1,1,0)$
3. $v = (0,0,1)$ and $v' = (1,1,0)$

are not realizable. We will check them one by one. Recall that by the Rosenstiehl Criterion the even edges form a cut in the interlacement graph if it is realizable. In Figure 3.19 the respective chord diagrams and interlacement graphs are shown and, evidently, the even edges do not form a cut in any of the three cases.

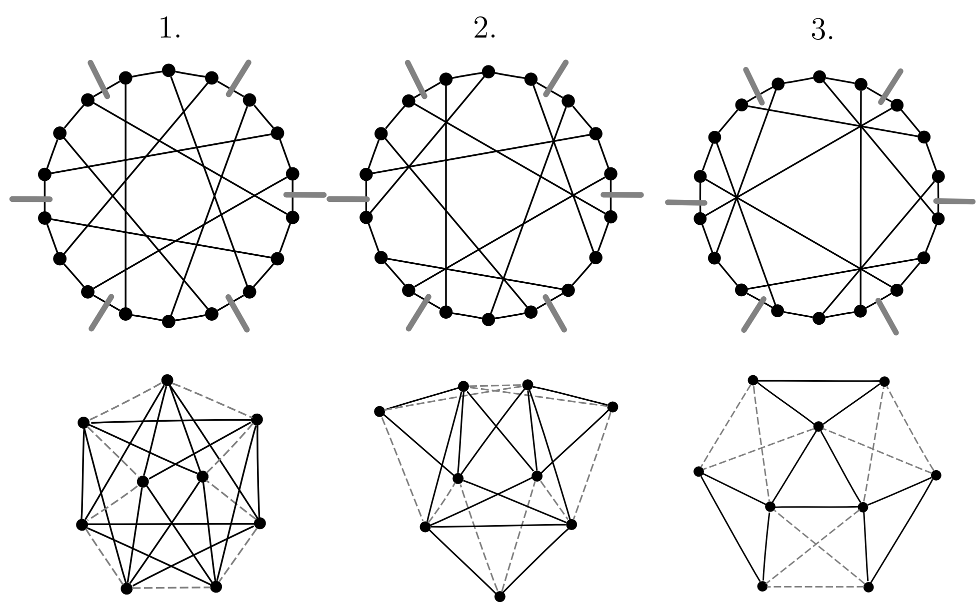

Figure 3.19. The chord diagrams and interlacement graphs of the three cases that remain to be checked. In the figure the even edges in the interlacement graphs are dashed.

□

3.6 Touching Thrackles

Given our approach to Gauss codes, the concept of a touching thrackle suggests itself. As the characterization of touching curves was much easier than the characterization of crossing curves, there is hope that the bound on the number of edges is much easier to show for touching thrackles than it is for (crossing) thrackles. And indeed this is the case.

A **touching thrackle** is a drawing of a graph G in the plane with the property that every two edges have exactly one point in common: either a common endpoint or a common interior point at which they *touch* each other. Martin Aigner asked whether graphs that have a touching thrackle drawing necessarily have at most as many edges as vertices. As C_3 and C_4 have a touching thrackle drawing (see Figure 3.20) this bound is best possible.

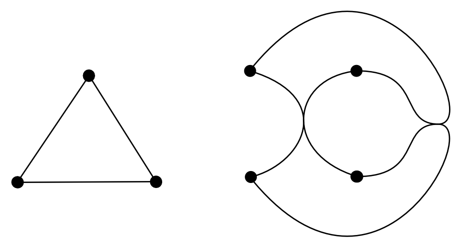

Figure 3.20. Touching thrackles drawings of the C_3 and the C_4.

Recall that the C_4 is the only circle that does not have a (crossing) thrackle drawing. Indeed, it turns out that regarding circles the situation with touching thrackles is the exact opposite of the situation with crossing thrackles.

3.14. All circles C_k for $k \geqslant 5$ do not have a touching thrackle drawing.

This is an immediate corollary of the following lemma.

3.15. If a graph G has a touching thrackle drawing and it contains an open or closed path P of 4 edges as a subgraph, then G contains no other edge.

Proof. Let $P = x_1\, a x_2\, b x_3\, c x_4 d x_5$, where possibly $x_1 = x_5$. Depending on whether or not $x_1 = x_5$, the drawing of P will be an open curve with 3 double points or a closed curve with 2 double points. In either case it will have 4 faces. We claim that the sets of edges that have points on the boundary of these faces are

$$\{a, b, c\}, \{a, b, d\}, \{a, c, d\}, \{b, c, d\}.$$

See Figure 3.21 for two examples. Note, that the statement of our claim is not entirely precise. We consider the face f' in Figure 3.21 to correspond to the set $\{a, b, c\}$ and yet d does have precisely 2 points on the boundary of f'. Hence, we refine our definition by specifying that an edge e is not in "the set" of a face f, if it has only finitely many interior points on the boundary of f. Note that if an edge e has an endpoint on the boundary of a face f, it automatically has infinitely many points on the boundary of f, because we are dealing with a path and hence every vertex has degree $\leqslant 2$.

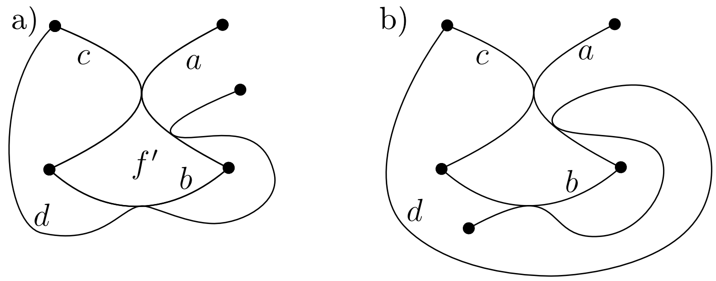

Figure 3.21. Two examples of touching thrackle drawings of an open path of 4 edges.

To see that our claim is true we argue as follows. By the thrackle condition, a and c touch at a double point p_{ac}. The segment of the path between the two occurrences of p_{ac} forms a simple closed curve that has two faces f_1 and f_2. The edge d has to be contained in one of these faces (except for finitely many points), as it cannot cross another edge. Let f_1 be the face d is contained in. Therefore, the edges that have infinitely many points on the boundary of f_2 are a, b and c. In the drawing of P only the curve segment $x_1 \overset{a}{\text{---}} p_{ac}$ can be contained in f_2 (and that only if P is open). The removal of this segment from f_2 does not separate f_2 and it does not change the set of curves that have infinitely many points on its boundary. Thus, $\{a, b, c\}$ is one of the sets we are looking for and, by symmetry, so is $\{b, c, d\}$.

Now, edge d has to touch b at a point p_{bd} and d has to have a point p_{ad} in common with a which may possibly be an endpoint. The segments

$$p_{ad} \overset{a}{\text{---}} x_2 \overset{b}{\text{---}} p_{bd} \overset{d}{\text{---}} p_{ad}$$

form a simple closed curve, with faces f_1' and f_2'. Let f_1' be the face not containing c. If P is closed, f_1' does not contain any other curve segments in our drawing of P, so f_1' is a face of P with corresponding set $\{a, b, d\}$. If P is open, f_1' may or may not contain the curve segments $x_1 \overset{a}{\text{---}} p_{ad}$ and/or $p \overset{d}{\text{---}} x_5$ where $p \in \{p_{ad}, p_{bd}\}$ is the double point d visits last. Removing these segments from f_1', however, does not separate f_1' or change the set of edges that have infinitely many points on its boundary, which is $\{a, b, d\}$. By symmetry there also is a face of P that has edges $\{a, c, d\}$ on its boundary. As P either is a closed curve with 2 double point or an open curve with 3 double points, its drawing has precisely 4 faces. Hence all faces are accounted for and we are done with the proof of our claim.

Now, suppose there is another edge e in the touching thrackle drawing of G. On the one hand e has to have a point in common with each of a, b, c and d, while on the other hand it has to be contained in a face of P. However, no face has all four edges on its boundary.

There is one fine point still to be considered: What about the finitely many interior points an edge may have on the boundary of a face? e is not allowed to contain any of these, as the thrackle condition forbids three edges to have a common interior point. Here we should note that if a vertex x is on the boundary of a face f, then all edges of P incident with x are contained in "the set" of f simply because they automatically have infinitely many points on the boundary of f. □

This lemma shows that there are only very few graphs that have a touching thrackle drawing. Still, there are infinitely many graphs that have a touching thrackle drawing and as many edges as vertices. These graphs can be constructed by picking any star and adding any one edge (see Figure 3.22).

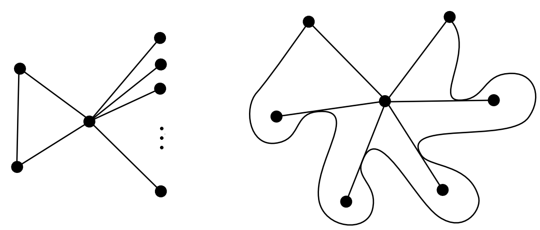

Figure 3.22. All stars with one additional edge (shown on the left) have a touching thrackle drawing. Such a drawing of the 6-star with one additional edge is shown on the right.

Given 3.15 it is easy to prove the bound on the number of edges for graphs that have a touching thrackle drawing.

3.16. Touching Thrackle Theorem

If a graph G has a touching thrackle drawing, then $|E(G)| \leqslant |V(G)|$.

Proof. A counterexample G cannot contain a C_4 as a subgraph, as G would contain one additional edge, which is impossible by 3.15. Now, as in Woodall's reduction of the Thrackle Conjecture, every counterexample G will have a component that contains two cycles, that is two subgraphs C_3 that do not have an edge in common. (If the two copies of C_3 shared an edge, a C_4 would result.) As they are in one component, there is a path of 0 or more edges linking the two. But then we find an open path of 4 edges in G, plus one additional edge, which again yields a contradiction. □

Now, what if we consider thrackles in which the edges may cross or touch each other? A graph that has such a drawing can have more edges than vertices as the following example of a C_4 with a single chord shows (Figure 3.23).

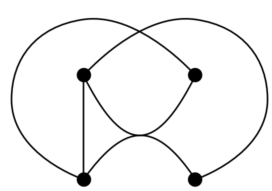

Figure 3.23. The K_4 minus one edge, drawn such that any two edges either have a common endpoint, or a common interior point.

So, what is the correct bound on the number of edges in this setting?

Bibliography

[1] M. Aigner. *What is a Map?* unpublished.

[2] G. Cairns and Y. Nikolayevsky. Bounds for generalized thrackles. *Discrete Computational Geometry*, 23:191–206, 2000.

[3] H. de Fraysseix and P. Ossona de Mendez. A short proof of a gauss problem. In G. D. Battista, editor, *Graph Drawing, Rome, Italy, September 18-20, 1997*, pages 230–235. Springer, 1998.

[4] T. Fleming and B. Mellor. Gauss codes for graphs. *arXiv:math.CO/0508269*, 2005.

[5] C. Godsil and G. Royle. *Algebraic Graph Theory*. Springer, New York, 2001.

[6] L. Lovász and M. Marx. A forbidden substructure characterization of gauss codes. *Bulletin of the American Mathematical Society*, 82:121–122, 1976.

[7] L. Lovász, J. Pach, and M. Szegedy. On conway's thrackle conjecture. *Discrete Computational Geometry*, 18:369–376, 1997.

[8] B. Mohar and C. Thomassen. *Graphs and Surfaces*. John Hopkins University Press, Baltimore and London, 2001.

[9] P. Rosenstiehl and R. C. Read. On the principal edge tripartition of a graph. *Annals of Discrete Mathematics*, 3:195–226, 1978.

[10] D. R. Woodall. Thrackles and deadlock. In D. J. A. Welsh, editor, *Combinatorial Mathematics and its Applications*, Proc. 1969 Oxford Combinatorial Conference, pages 335–347, 1971.