Rigidity of thin disk configurations

by

Andrey Mikhaylovich Mishchenko

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)
in the University of Michigan
2012

Doctoral Committee:

Professor Jeffrey C. Lagarias, Chair
Professor Daniel Jacobson
Professor G. Peter Scott
Professor Mario Bonk, University of California, Los Angeles
Professor Kai Rajala, University of Jyväskylä

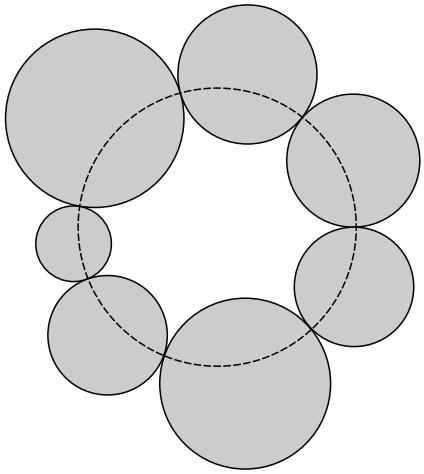

To Karen, Victor, and Claudia

Acknowledgments

Thanks Jeff for taking me on as a student, who knows why you actually agreed to it, and for being the first one to ever mention circle packings to me. Also I suppose you might think I'm joking but it always cheers me up to see you!

Thanks Mario, Sergiy, and Steffen, for kind-of being surrogate advisors. Special thanks go out to Mario for giving me a C in first-year complex analysis!

Thanks NSF grants DMS-0456940, DMS-0555750, DMS-0801029, DMS-1101373 for many fruitful semesters off from teaching.

Thanks to my undergraduate teachers, especially for writing me letters that somehow got me admitted to Michigan. Thanks Juha for admitting me.

Thanks dad for nurturing my math brain when I was a kid, sorry I didn't end up a physicist. Thanks mom for nurturing my "romantic"–crazy side, sorry I didn't end up a karate yoga shaman. Thanks Natasha for all the mornings playing frisbee at the park, and thanks Sergei for being the responsible one in charge of the house now that I'm gone.

Thanks Peter Scott for teaching me that a picture is often the best form of proof. Thanks Dan Jacobson for kindly serving as a cognate member, and for (implicitly) challenging me to write something down that a non-mathematician could actually maybe read. Thanks Lars and Juan for many fruitful discussions.

Thanks to my friends from undergrad, who got watch me try to be social for the first time. Even though you'll never read this.

Thanks to my grad student friends from grad school, who got to watch me rise to my own personal ultimate pinnacles of social grace. Special thanks to Jordan, Julian, and Michael for the philosophical arguments. Special thanks to Ross and Jessie for the trip to Mexico, my first time really out of the country. Special thanks to Jay and Silent Bob for the first-year trip to NY, especially the pit stop at Perkins. Special thanks to Nina, is everything alright? Special thanks to the current second-year class for all the table tennis games. I'd like to especially thank Chris and Wendy for a nice enjoyable skiing trip with the fellowship of the ring, and for educating me on hipster culture and Japanese culture, respectively. Special thanks to Max, for

somehow always ending up as my teammate in puzzle challenges, bridge games, and frisbee tournaments. Thanks to the whole after school math club. Hello Hunter how are you doing? Also thanks guys for helping me take my time writing this thesis.

Special thanks to Kanwar for being the cleanest and tidiest roommate anyone could ever ask for. Special thanks to Patrick for being the best Scottish roommate anyone could ever ask for.

Special thanks to Zach the Mack

Thanks Fidel and Gourab for the New England trip, which I will never forget and I'm pretty sure you guys are still trying to forget.

Thanks to Victor for being such a nice, humble, sweet, soft-spoken, gentlemanly individual, and always an excellent influence.

Thanks Karen for being the only friend I have ever had who never had any drama (with me).

Thanks Kai for always lighting up the room with your smile!

Thanks to the people of Finland for their hospitality.

Thanks Helena for being the best customer a cookie salesman could ever ask for. Thanks Tapio for always being a joy to some of the senses. Thanks Sanelma for always being ready to see who is better at something, snowboarding, not talking, etc.

Thanks Claudia for being such a cute little puppy. Also thanks for putting up with me.

Actually, thanks everyone for putting up with me!

truck.

Preface

To the non-mathematician: Chapter 1 is intended to be readable by anyone, without any background whatsoever. By the end of Chapter 1, we state the main result of this thesis. For the reader not accustomed to reading mathematics, we give some advice at the end of this preface.

The majority of this thesis is self-contained given a familiarity with standard classical complex analysis, on the level of [Ahl78] or [Rud87]. The only significant exceptions are some parts of the introduction to circle packing, which lean on the fundamentals of hyperbolic geometry and Riemann surface theory. For an introduction to hyperbolic geometry, see [BP92], and for a reference text see [Rat06]. For a text on the theory of Riemann surfaces, see [For91]. Stephenson gives brief primers on both of these areas in his introductory book on circle packing, [Ste05].

This thesis is structured as follows:

- Chapter 1 is an introduction to the thesis intended for non-mathematicians. It also serves as a gentle introduction for mathematicians unfamiliar with circle packing. We make two remarks: (1) for ease of exposition, the definition of *thin* given in Chapter 1 is slightly stronger than the precise technical definition given later, so the statement of the main result given in this chapter is slightly weaker than what we actually prove, and (2) the order in which the material is presented is slightly different between Chapters 1 and 2, also for ease of exposition.

- Chapter 2 introduces circle packing. We state some major theorems from the area and discuss applications. We then introduce disk configurations, which generalize the notion of circle packings, and state our main results and conjectures. We end with some expository sections and a list of further references.

- Chapter 3 introduces our main tool, fixed-point index. We state and prove some fundamental lemmas on fixed-point index. At the end, we state our main technical result, which we call the Index Theorem, and a conjectured generalization.

- In Chapter 4, we first sketch the proof of the Index Theorem. We then discuss some key geometric lemmas and propositions. Finally, we give the proof of our main results using the machinery just developed.

- Chapters 5 and 6 are spent proving the Index Theorem. In particular, Chapter 5 consists of elementary topological and geometric arguments. In Chapter 6, we develop a tool for working with fixed-point index, the so-called torus parametrization, and then apply this tool to work out the remaining details.

This thesis is typeset in LaTeX, using many $\mathcal{A}\mathcal{M}\mathcal{S}$-LaTeX macros, among others. Most of the figures were drawn in LaTeXDraw, available at `http://latexdraw.sourceforge.net/`, and are rendered using `pstricks`, often after some code tweaking. Some figures were hand-written in `pstricks` code. Incidentally, I highly recommend LaTeXDraw for generating figures. The two schematics on p. 41 were typeset using the commutative diagram package XY-pic.

The frontispiece is a cycle of tangent disks with a dual circle. The idea to study dualizable circle packings ultimately led to the results of this thesis.

Advice on reading mathematics: Chapter 1 is intended to be readable without any background. However, reading math is different from reading most other kinds of writing, and can be difficult and frustrating, so here are some tips to make it easier:

- Whenever a concept is introduced for the first time, draw some examples for yourself. This will make it easier to internalize the concept.

- If something is stated as a "fact," then either we have argued in support of it in the text, or it should be possible to convince yourself that it is true by drawing a few examples. If something is stated as a "theorem," then it is true, but it is difficult to see why it is true, and difficult to prove it.

- Go slowly. A page in a novel will rarely take more than a minute or two to read, but reading a single page of unfamiliar math can take much longer.

- Do the exercises. They will help you to follow along. Actually doing the exercises organizes the new ideas in your mind, and will make reading on much easier and ultimately faster than if you didn't do the exercises at all.

Contents

List of Figures

Abstract

The main result of this thesis is a rigidity theorem for configurations of closed disks in the plane. More precisely, fix two collections $\mathcal{C}$ and $\tilde{\mathcal{C}}$ of closed disks, sharing a contact graph which (mostly-)triangulates the complex plane, so that for all corresponding pairs of intersecting disks $D_i, D_j \in \mathcal{C}$ and $\tilde{D}_i, \tilde{D}_j \in \tilde{\mathcal{C}}$ we have that $\measuredangle(D_i, D_j) = \measuredangle(\tilde{D}_i, \tilde{D}_j)$. We require the extra condition that the collections are *thin*, meaning that no pair of disks of $\mathcal{C}$ meet in the interior of a third, and similarly for $\tilde{\mathcal{C}}$. Then $\mathcal{C}$ and $\tilde{\mathcal{C}}$ differ by a Euclidean similarity.

Our proof is elementary, using essentially only plane topology arguments and manipulations by Möbius transformations. In particular, we generalize an argument which was previously used by Schramm to prove the rigidity of configurations of pairwise interiorwise disjoint closed disks having contact graphs triangulating the complex plane. It was previously thought that his proofs depended too crucially on the pairwise interiorwise disjointness of the disks for there to be a hope for generalizing them to the setting of configurations of overlapping disks.

Analogous versions of our rigidity theorem have easier proofs via a discrete version of the Maximum Modulus Principle in the case where $\mathcal{C}$ and $\tilde{\mathcal{C}}$ share a contact graph which (mostly-)triangulates the hyperbolic plane, or the Riemann sphere. We describe these proofs as well. These are relatively straightforward generalizations of the corresponding proofs in the case of configurations of pairwise interiorwise disjoint disks. Then by a simple argument via covering space theory and the Uniformization Theorem, we get an analogous rigidity statement for thin disk configurations having contact graphs triangulating an arbitrary Riemann surface.

We include a brief and gentle introduction intended for non-mathematicians. Then we give a survey of the field of circle packing, which is the area that our result fits into. We also state some open problems and conjectures from this area, including conjectured generalizations both of our main result and of our main technical theorem.

Chapter 1

Introduction for non-mathematicians

Roughly speaking, this thesis studies the patterns that circles can make with one another. This area of study is known as *circle packing*. For the first half of this introduction, we explore some natural, classical questions about circle packing.

In this thesis we consider configurations of disks more general than classical circle packings. In particular, we allow our disks to overlap, whereas the disks of a circle packing are required not to overlap. We can ask questions about our disk configurations analogous to those we asked about circle packings. The answers to some of these questions make up the main results of this thesis. We describe these questions and answers for the second half of this introduction.

1.1 What is a circle packing?

Suppose we arrange some coins flat on a table. We get a picture that looks something like this:

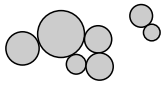

Figure 1.1: **Our first example of a circle packing.**

The coins may be of different sizes. Also, the coins may touch, but for now we insist that they do not overlap. Configurations such as these are traditionally known as *circle packings*:

Definition. A ***circle packing*** is a collection of disks which don't overlap.

Roughly speaking, we will be exploring the patterns that can be formed by the disks in a circle packing. For example, we say that somehow, intuitively, the following two circle packings have "the same pattern":

(a) (b)

Figure 1.2: **Two circle packings which have "the same pattern,"** in some sense.

We wish to make the notion of the "pattern" of the disks of a circle packing more precise. In order to do so, we first introduce the related notion of a *graph*:

Definition. A ***graph*** is a collection of ***vertices***, which we usually draw as dots, and ***edges***, which are connections between the dots.

The easiest way to represent a graph is via a drawing. However, we stress that a graph is an abstract object which is independent of any drawn representation of it.

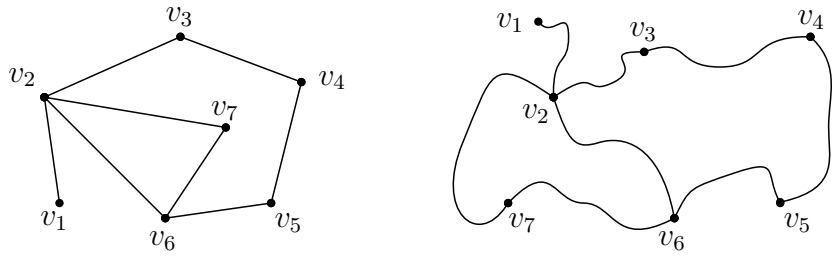

(a) (b)

Figure 1.3: **Two different representations of the same graph.** The graph we have drawn here has 7 vertices. Here the vertices have been labeled $v_1, v_2, \ldots, v_7$. Then there is an edge between v_i and v_j in (a) whenever there is an edge between v_i and v_j in (b) and vice versa. Note also that for example, we do not consider that there is an edge between v_3 and v_5 in this graph, because we cannot get from v_3 to v_5 without passing through other vertices.

Graphs are all around us. For example, the *handshake graph* is the graph having a vertex for every person, and having an edge between two people if they have ever shaken hands. You may be more used to hearing graphs referred to as *networks*. For example, the *Facebook friend graph* is the graph having a vertex for every Facebook account, and having an edge whenever the two accounts are friended on Facebook. Of course, the graphs in these two examples both have huge numbers of vertices, and it would be very hard to draw representations of them. We suggest an easier exercise:

Exercise 1. Draw a representation of the following graph:

- The vertices are the English words for the whole numbers from one to six.
- Two vertices have an edge connecting them whenever the two corresponding words share a letter.

The concept of a graph helps us formalize what we mean when we discuss the "pattern" formed by the disks of a circle packing. In particular, we may use a graph to capture and isolate the information of "how the disks of the packing meet":

Definition. Suppose we have a circle packing $\mathcal{P}$ consisting of the disks $D_1, D_2, D_3, \ldots$. Then the ***contact graph*** of the packing $\mathcal{P}$ is the graph described as follows:

- The vertices are the disks $D_1, D_2, D_3, \ldots$ of the packing.
- Two vertices have an edge connecting them if and only if the corresponding disks touch.

We explore two examples:

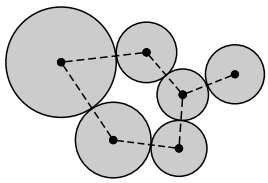

Figure 1.4: **A circle packing and its contact graph.** The edges of the graph are drawn as dashed lines. This is the easiest way to visualize the contact graph of a circle packing, and usually the easiest way to draw a representation of it.

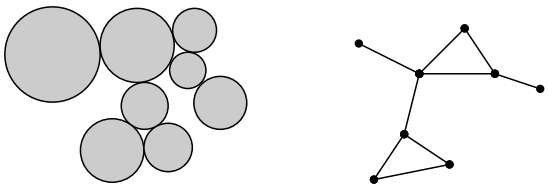

(a) (b)

Figure 1.5: **Another circle packing and its contact graph.** Here we have drawn the contact graph separately from the packing, to again emphasize that the graph is an abstract object distinct from any drawn representation of it.

Exercise 2. Draw the contact graph of the following circle packing:

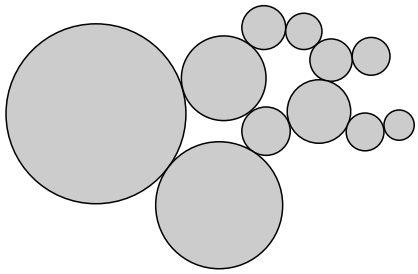

Experience leads us to conclude:

Observation. *If we start with a circle packing, it is not so hard to draw its contact graph.*

In other words, it is easy to "go from" a circle packing to its contact graph. Then we might ask, when can we go "in the other direction?" This is our segue into the next section.

1.2 Existence of circle packings

To build intuition, we start this section with an exercise:

Exercise 3. Draw circle packings having the following contact graphs:

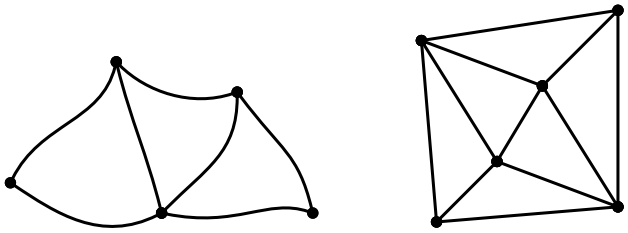

This leads us to consider the following natural question:

Existence Question. *If we start with a graph G, is there always a circle packing having G as its contact graph?*

For simplicity, in our study of the existence question, we will consider only *finite* graphs, meaning graphs which have finitely many vertices and edges. We first observe that circle packings' contact graphs never have *loops*, nor *repeated edges*:

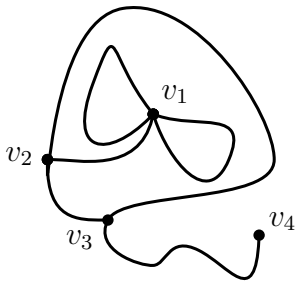

Figure 1.6: **A graph which *does* have loops and repeated edges.** The vertices v_2 and v_3 have two edges connecting them, an example of what is meant by a *repeated edge*. The vertex v_1 has two loops on it: a *loop* is an edge from a vertex to itself.

A graph without loops and repeated edges is called *simple*. Thus we have our first necessary condition for a graph to appear as the contact graph of a circle packing:

Existence Condition 1. *If G is the contact graph of a circle packing, then G is simple.*

There is another, more subtle observation we need to make to answer our existence question for circle packings:

Fact. *Not every graph can be drawn without its edges crossing each other.*

The idea of *edges crossing* is illustrated in the following example:

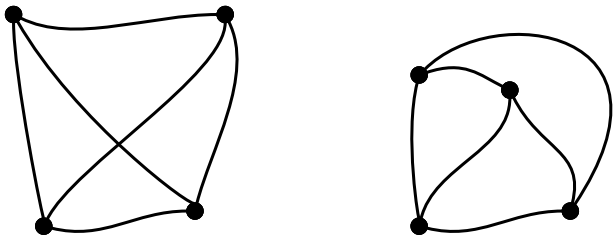

Figure 1.7: **Two representations of the same graph, one planar and one not.** The graph represented here has four vertices, and the vertices are drawn bigger than usual for emphasis. In (a), two edges of the graph cross. This does not happen in (b). The drawing in (b) is called a *planar representation* for this graph.

A graph is called *planar* if it has a planar representation, meaning that it can be drawn (in the plane) without edge crossings. The source of the word *planar* is that the space we are working in, where geometry is usually done, is called the Euclidean plane, or just the *plane*. To internalize the idea of planarity and the fact that not all graphs are planar, we give our next exercise:

Exercise 4. Try to draw a representation of the following graph, so that the edges don't cross each other:

- It has five vertices, and
- *every pair* of vertices has an edge connecting them.

In other words, draw five dots on a piece of paper, and try to draw a connection between every pair of the dots so that these connections don't cross. (It is impossible, but try to convince yourself of this, rather than taking it on faith.)

On the other hand, suppose that the graph G is the contact graph of some circle packing. Then we may draw a representation of G as we did in Figure 1.4, by putting the vertex corresponding to a disk D at the center of D, and connecting the vertices of touching disks by straight line segments. This representation of G will be planar because the disks of our packing do not overlap. Thus we have found our second necessary condition for a graph to appear as the contact graph of a circle packing:

Existence Condition 2. *If G is the contact graph of a circle packing, then G is planar.*

Thus for example, it is impossible to find a circle packing whose graph is the one described in Exercise 4.

We summarize our findings so far:

Fact. *If G is the contact graph of a circle packing, then G is simple, meaning that it has no loops nor repeated edges, and planar, meaning that it can be drawn (in the plane) without edge crossings.*

This is just the combination of conditions 1 and 2. The amazing thing is that these conditions are the *only ones* we need for a positive answer to our existence question:

Existence Theorem. *If G is a simple planar graph, then there is some circle packing having G as its contact graph.*

The Existence Question for circle packings was first asked by Paul Koebe, and he answered it himself in 1936, by proving the Existence Theorem we just stated. The proof uses *complex analysis*, the theory of calculus on the complex numbers $\mathbb{C} = \{a + bi\}$. It turns out that the field of circle packing is a kind of discrete analog of complex analysis. This is a major motivating factor for the study of circle packings.

1.3 Rigidity of circle packings

We have seen that different circle packings can share a common contact graph. For example, the packings in Figure 1.2 on p. 2 have the same contact graph. In this section, we explore the question of whether there is a contact graph so that *only one* circle packing has that contact graph.

Of course, if we start with a circle packing, pick the whole packing up as a single unit, move it around, and then set it back down, we will get a new circle packing having the same contact graph. Thus it will never really be the case that there is only one circle packing having some given contact graph. If we are to hope for a uniqueness statement, we need a way around this problem. This leads us to explore the notion of *similar packings*.

Recall from elementary school geometry that two triangles are called *similar* if one is a copy of the first, but scaled up or down and moved around:

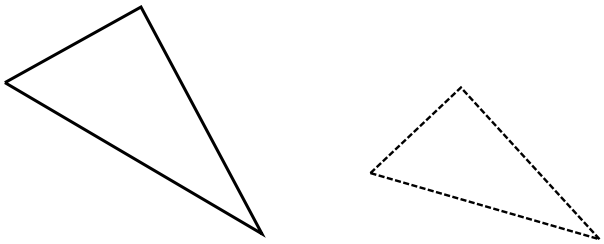

Figure 1.8: **Two similar triangles.**

We can define two packings to be similar in an analogous way:

Definition. Two circle packings are called ***similar*** if it is possible to get from one to the other by combining the following four operations:

- sliding,
- rotating,
- reflecting, and
- scaling.

We stress that it is essential that the operations are applied to the *whole packing at once*, not to the disks one at a time. If we were allowed to apply the operations to one disk at a time, then it would be possible to get from a packing to any other one having the same number of disks via these operations, which is not what we want.

If two packings are similar, then we consider them to be "essentially the same":

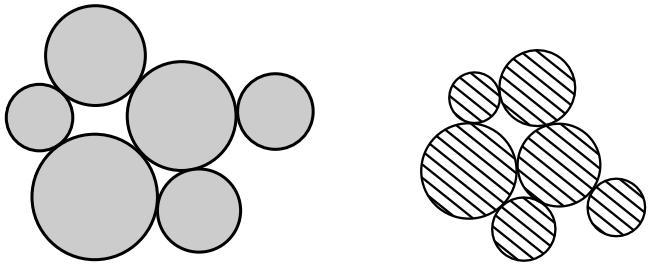

Figure 1.9: **Two circle packings which are "essentially the same."**

The reason for this is that similar packings *always* have the same contact graph. On the other hand, there may be some circle packing $\mathcal{P}$ so that *every* packing which has the same contact graph as $\mathcal{P}$ turns out to be similar to $\mathcal{P}$. Thus we ask the following question:

Uniqueness Question. *When is there an essentially unique circle packing corresponding to a given graph?*

For every circle packing we have seen so far, it is possible to find an essentially different packing, meaning a packing which is not similar to the original one, which has the same contact graph. In fact, it is not too hard to believe the following:

Fact. *If a circle packing $\mathcal{P}$ has finitely many disks, then there are circle packings having the same contact graph as $\mathcal{P}$, but which are not similar to $\mathcal{P}$.*

Thus, we need to turn our attention to *infinite packings*, circle packings having an infinite number of disks:

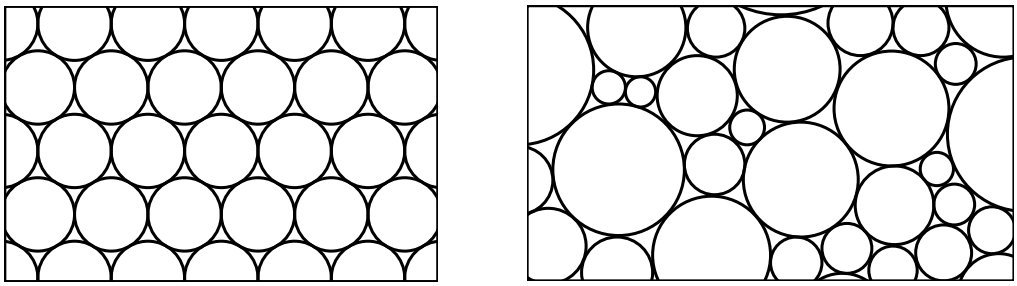

(a) (b)

Figure 1.10: **Pieces of infinite circle packings.** Of course, it is impossible to draw the entirety of a packing that goes on forever. However, sometimes, it is still possible to completely describe the packing. For example, the packing shown in (a) is the infinite packing by disks which all have the same size, so that each disk has exactly 6 neighbors. This is known as the "penny packing." We leave what happens in (b) to the imagination.

We might be tempted to believe that if a packing "goes on forever in all directions," then it is essentially the only packing having its contact graph. Unfortunately, this does not turn out to be true. For example:

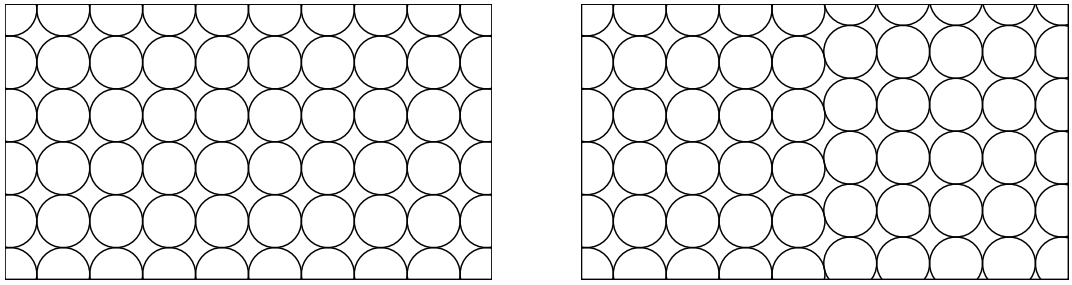

(a) (b)

Figure 1.11: **Two infinite packings which are not similar, but which have the same contact graph.**

Exercise 5. Describe the contact graph of the packings in Figure 1.11. Can you find a third packing which has the same contact graph, but which is not similar to either of (a) and (b)?

In Figure 1.11, we "perturbed" some disks of the packing, leaving others where they were, to get a new, fundamentally different packing having the same contact graph. On the other hand, it does not seem so easy to perturb in a similar way the disks of either packing in Figure 1.10. This is because of a special property of the contact graph:

Definition. Suppose that $\mathcal{P}$ is an infinite circle packing. Represent its contact graph G by putting a vertex at every circle's center, and connecting vertices of touching circles by straight segments. Suppose that then the graph G "cuts the plane into triangles." In that case we say that the packing $\mathcal{P}$ is ***triangulated***.

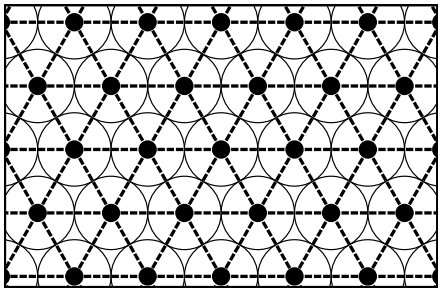

Figure 1.12: **The contact graph of the penny packing.** The key point is that every point of the plane lies inside one of the triangles formed by the contact graph. Thus the penny packing is *triangulated*.

Exercise 6. Draw the contact graph on the packing in Figure 1.10b, as we did in Figure 1.12 for the packing in Figure 1.10a. Verify that the packing in Figure 1.10b appears to be triangulated. (Of course, we don't really know what happens outside of the rectangular region we have shown.)

The following amazing fact turns out to be true:

Uniqueness Theorem. *If two circle packings have the same contact graph and are are both triangulated, then they are similar. In other words, if a packing $\mathcal{P}$ is triangulated and has contact graph G, then $\mathcal{P}$ is essentially the* only *circle packing having contact graph G.*

A alternative way to phrase this theorem is to say that triangulated circle packings are *rigid*, meaning we cannot modify a triangulated circle packing without changing its contact graph, except by applying rotations, scaling, etc. Compare this to the situation in Figure 1.11, where it is actually possible to modify the packings significantly while keeping the contact graph the same.

Furthermore, considering the example given in Figure 1.11, the following is not hard to believe:

Fact. *Suppose a circle packing $\mathcal{P}$ is not triangulated. Then there are circle packings having the same contact graph as $\mathcal{P}$, but which are not similar to $\mathcal{P}$.*

Thus the condition that $\mathcal{P}$ is triangulated is *necessary* for $\mathcal{P}$ to be rigid. The Uniqueness Theorem says that this condition is also *sufficient*. Thus the uniqueness guaranteed by the Uniqueness Theorem is the best that we could possibly hope for.

As we mentioned earlier, Koebe initiated the study of circle packings in 1936, and proved the Existence Theorem for circle packings. He also managed to prove a different kind of uniqueness theorem for circle packings: instead of considering packings in the plane, as we are, he considered packings on the surface of a sphere. The truth or falsehood of the Uniqueness Theorem we have stated here was unknown for a long time after Koebe. A proof was finally found in 1991 by Oded Schramm.

1.4 Overlapping disks

The results of this thesis stem from the following simple, open-ended question:

Question. *What if we allow the disks to overlap?*

To give some answers to this question, we will study so-called *disk configurations*:

Definition. A ***disk configuration*** is any collection of disks so that no disk in the collection is completely contained inside of another.

As usual, we start with an example:

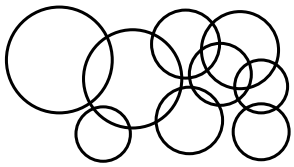

Figure 1.13: **Our first example of a disk configuration** (which is not also a circle packing)

When discussing disk configurations, it will be helpful to have a notion of *overlap angle*:

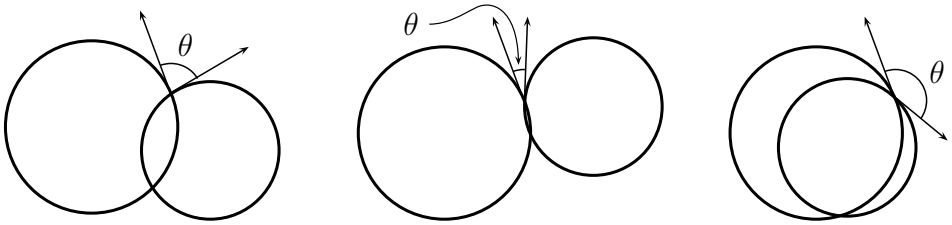

Figure 1.14: **The overlap angles between several pairs of disks.** In each case, we have chosen a point where the boundary circles meet, and drawn the tangent rays to the circles pointing out of the disks. Then the overlap angle θ is the angle between these rays. We note two facts: (1) the overlap angle gets bigger as the disks "move closer together," or gets smaller as the disks "move apart," and (2) the overlap angle is the same at both points of intersection of the boundary circles.

The *contact graph* of a disk configuration is defined the same way as the *contact graph* of a circle packing: it is the graph with a vertex for every disk, so that two vertices are connected by an edge whenever the corresponding disks meet. In the case of disk configurations, however, we label each edge of the contact graph with the overlap angle between those two disks. For example:

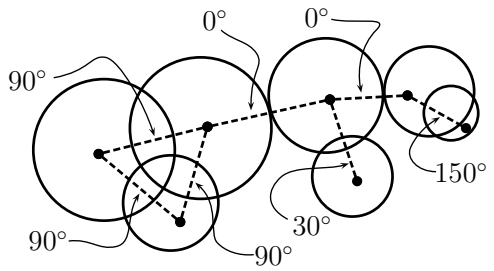

Figure 1.15: **The contact graph and overlap angles of this disk configuration.**

To internalize these concepts, we suggest our next exercise:

Exercise 7. Draw a disk configuration having the following contact graph and (approximate) overlap angles:

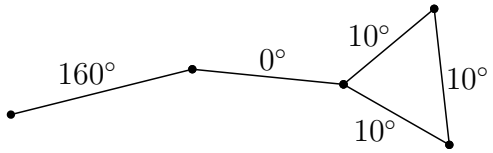

More generally, we ask ourselves the following natural question:

Question. *Given a graph G, so that its edges are labeled with prescribed overlap angles, does there exist a disk configuration having G as its contact graph with the required overlap angles?*

The answer to this question turns out to be very messy compared to the answer to our Existence Question for circle packings. For our first example, consider the following:

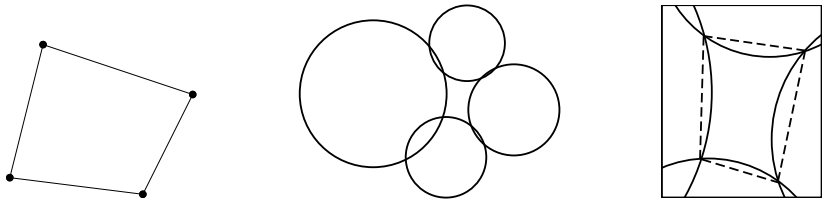

Figure 1.16: **Four disks in a closed chain.** Suppose four disks have the contact graph shown in (a). An example of four such disks is shown in (b). If we zoom in on the hole formed by the four disks, we see the picture in (c). The sum of the angles inside the dashed quadrilateral we have drawn is clearly bigger than the sum of the overlap angles of the four disks. On the other hand, the sum of the angles inside of a quadrilateral is always exactly 360°. So if four disks have the contact graph shown in (a), then the sum of their overlap angles *must* be less than 360°.

Thus for instance there is no disk configuration having the following contact graph and overlap angles:

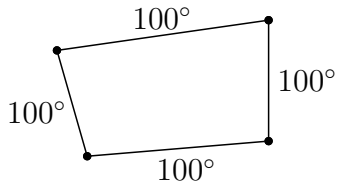

Figure 1.17: **An impossible contact graph given the indicated overlap angles**, impossible because the sum of the indicated overlap angles is greater than 360°.

This might lead us to conclude that no disk configuration has a contact graph containing, as a sub-graph, the graph pictured in Figure 1.17 including edge labels. However, this conclusion does not turn out to be true. Consider the following example:

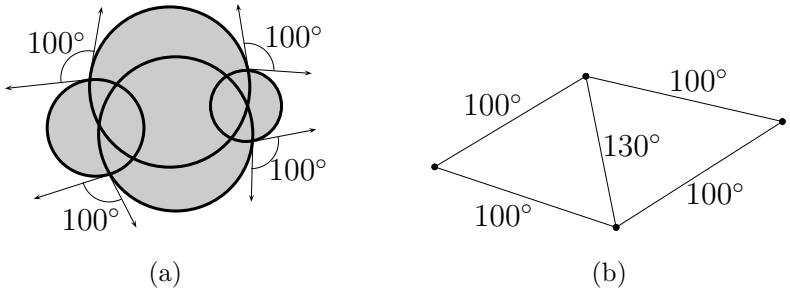

Figure 1.18: **Another disk configuration, and its contact graph with overlap angles labeled.**

At this point, we give up on the question of existence and summarize our findings thus far in the following observation:

Observation. *Given a graph G with edges labeled with prescribed overlap angles, it is not so easy to determine whether or not there is a disk configuration which has G as its contact graph, and has the required overlap angles.*

We reassure the reader that more satisfying answers are known, and have been since the 1970s, but they are rather technical to state here. Instead, we move on to the new result proved in this thesis:

1.5 Rigidity of thin disk configurations

The following is the disk configuration analog of the Existence Question we asked about circle packings:

Question. *Suppose we start with a disk configuration. Can we give conditions that guarantee that there is essentially no other disk configuration having the same contact graph and overlap angles?*

As before, the following is not too hard to believe:

Fact. *Suppose that the disk configuration $\mathcal{C}$ consists of only finitely many disks. Then there are disk configurations having the same contact graph and overlap angles as $\mathcal{C}$, but which are not similar to $\mathcal{C}$.*

We consider two disk configurations to be *significantly different* if they are not similar. What it means for two disk configurations to be similar is the same as what it meant for two circle packings to be similar.

Exercise 8. Draw a disk configuration which has the same contact graph and overlap angles as the configuration depicted in Figure 1.15, but which is not similar to it.

Also, as with circle packings, there are infinite disk configurations which have the same contact graph and overlap angles, but which are not similar:

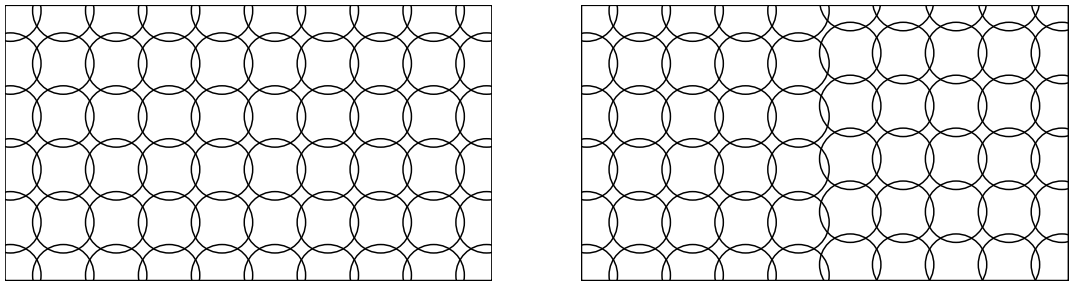

(a) (b)

Figure 1.19: **Two infinite disk configurations which are not similar, but which have the same contact graph and overlap angles.**

Similarly to before, it will help to consider *triangulated* disk configurations: a disk configuration is *triangulated* if its contact graph "cuts the plane into triangles." The main result of this thesis is the following:

Main Theorem. *If two thin triangulated disk configurations have the same contact graph and overlap angles, then the two disk configurations are similar.*

We still need to define the extra adjective *thin* in the statement of the Main Theorem:

Definition. A disk configuration is ***thin*** if no three of its disks meet at any point.

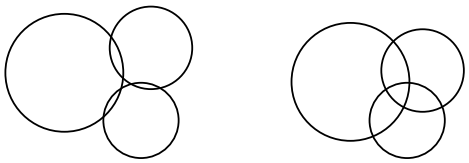

(a) (b)

Figure 1.20: **A disk configuration which is thin and one which is not.** The configuration in (b) is not thin, because there are points which are in all three of the disks at once. This does not happen in (a), so the configuration in (a) is thin.

We give two examples of disk configurations to which our Main Theorem applies:

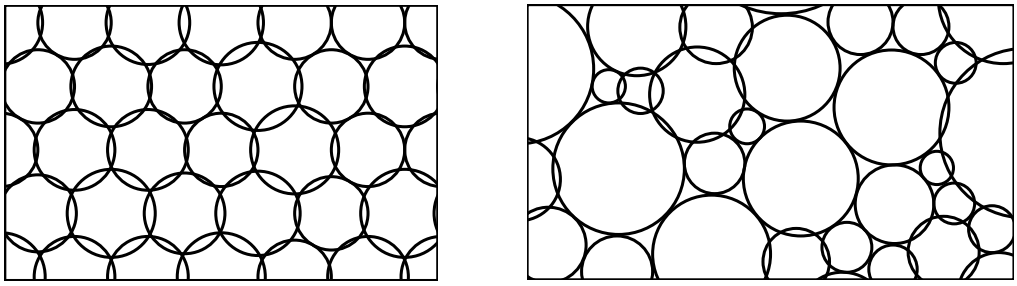

(a) (b)

Figure 1.21: **Pieces of thin triangulated disk configurations.**

We also give an example of a disk configuration which is not covered by our theorem:

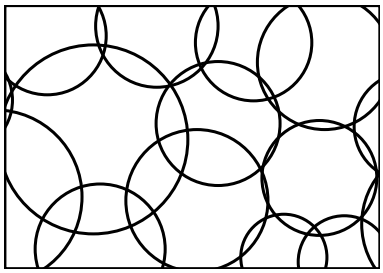

Figure 1.22: **A triangulated, but non-thin, disk configuration.**

Another way of stating our Main Theorem is to say that thin triangulated disk configurations are *rigid*, meaning that we cannot modify them while keeping the same contact graph and overlap angles, except by applying rotations, scaling, etc. Mathematicians refer to theorems of this type as *rigidity theorems*.

As before, consideration of Figure 1.19 makes the following believable:

Fact. *Suppose that $\mathcal{C}$ is a thin disk configuration which is not triangulated. Then there are disk configurations having the same contact graph and overlap angles as $\mathcal{C}$, but which are not similar to $\mathcal{C}$.*

Thus the uniqueness guaranteed by our Main Theorem is the best uniqueness we could hope for for thin disk configurations.

We use the thinness condition in our proof of the Main Theorem, but suspect that it is unnecessary, and make the following conjecture:

Conjecture. *Our Main Theorem is still true if we completely get rid of the thinness requirement from the statement.*

A conjecture is a mathematical statement that we believe to be true, but cannot prove. In 1999 Z.-X. He proved the following:

Theorem. *Our Main Theorem is still true if we get rid of the thinness requirement from the statement, provided that none of the overlap angles is bigger than* $90°$.

However, the restriction that overlap angles stay below $90°$ is a strong one, and he never published a version of this theorem where he managed to eliminate this restriction.

1.6 Why study this stuff?

An often-asked question upon learning the content of this thesis is, what real-world applications does it have? For insight into our motivations, we contemplate the following stanza by Keats:

> When old age shall this generation waste,
> Thou shalt remain, in midst of other woe
> Than ours, a friend to man, to whom thou sayst,
> "Beauty is truth, truth beauty," – that is all
> Ye know on earth, and all ye need to know.

Beyond this, ostensibly some people are trying to use circle packings to map regions of the brain, c.f. [Ste05, Section 23.4].

Chapter 2

Introduction to circle packing

2.1 Origins of circle packing

Traditionally, a *circle packing* is defined to be a finite collection of pairwise interiorwise disjoint metric closed disks in the Riemann sphere $\hat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$ equipped with the constant curvature $+1$ metric as usual. There are no conditions on the radii of the disks. Nowadays some authors use the phrase "circle packing" more loosely to refer to collections of disks which may overlap, but for clarity we do not do so in this thesis.

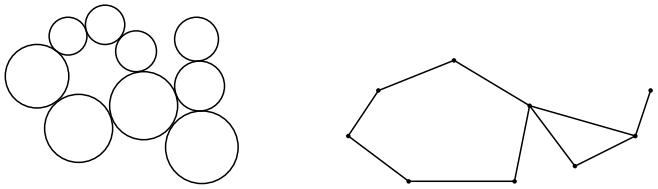

Figure 2.1: **A circle packing and its contact graph.** Although the circle packing gives us a natural geodesic embedding of its contact graph, we draw the graph separately here to emphasize that it is a purely combinatorial object.

We say that two closed disks are *tangent* if their boundary circles are tangent. Let $\mathcal{P}$ be a circle packing on the Riemann sphere. The *contact graph* of $\mathcal{P}$ is the graph G whose vertex set is in bijection with the disks of $\mathcal{P}$, so that $\langle v_i, v_j \rangle$ is an edge in G if and only if D_i and D_j are tangent. We adopt the convention that $D_i \in \mathcal{P}$ is the disk corresponding to $v_i \in V$ where (V, E) is the contact graph of $\mathcal{P}$. We will also sometimes use the notation that $D_v \in \mathcal{P}$ is the disk corresponding to $v \in V$, when it is clear what is meant. The following is a natural question:

Question 2.1. *Given a graph G, does it occur as the contact graph of some circle packing?*

A circle packing gives us a natural embedding of its contact graph, so if G is the contact graph of a circle packing on the Riemann sphere then G must be planar. It

turns out that this is essentially the only requirement on G for a positive answer to Question 2.1:

Circle Packing Theorem 2.2. *Every simple planar graph G occurs as the contact graph of some circle packing.*

Recall that a graph is called *simple* if it is undirected, does not have loops, and has no repeated edges. A *loop* is an edge from a vertex to itself. A graph is called *planar* if it can be drawn in the plane with no edge crossings.

We will sketch a proof of the CPT 2.2 in the case where G is finite. The construction we give is due to Koebe, and appears in [Koe36]. It relies on a special case of the so-called Koebe Conjecture, posed in [Koe08, p. 358]:

Koebe Conjecture 2.3. *Every domain $\Omega \subset \hat{\mathbb{C}}$ is biholomorphically equivalent to some circle domain $\circledcirc$. Furthermore both $\circledcirc$ and the biholomorphism $\varphi : \Omega \to \circledcirc$ are uniquely determined by Ω, up to post-composition of φ by Möbius transformations.*

A *domain* in $\hat{\mathbb{C}}$ is a connected open subset of $\hat{\mathbb{C}}$, and a *circle domain* is a domain so that every boundary component is a metric circle or a point. Thus the KC 2.3 is a conjectured generalization of the classical Riemann Mapping Theorem (see [Ahl78, Chapter 6, Theorem 1; Rud87, Theorem 14.8]):

Riemann Mapping Theorem 2.4. *Any simply connected domain $\Omega \subsetneq \mathbb{C}$ is biholomorphically equivalent to the open unit disk $\mathbb{D}$, and the biholomorphism is uniquely determined by Ω, up to post-composition by Möbius transformations $\mathbb{D} \to \mathbb{D}$.*

Koebe himself stated the following, c.f. [Koe08]:

Koebe's Theorem 2.5. *The Koebe Conjecture 2.3 holds for domains having finitely many boundary components.*

His construction of the uniformizing map is via an iterative algorithm, and goes as follows. Let $\Omega = \Omega_1$ denote our domain, and number its boundary components $1, \ldots, n$. Let Ω_{i+1} be obtained from Ω_i by applying a biholomorphism sending the $(i \mod n)^{\text{th}}$ boundary component γ to a metric circle. We may do this by applying the Riemann Mapping Theorem 2.4 to the domain in the Riemann sphere which is bounded by γ and which contains Ω_i. Of course, after the next iteration of this algorithm, probably that boundary component will have been distorted into a Jordan curve which is not a circle, but this new distortion is "not too bad." Anyway, one shows that under the appropriate normalizations, the Ω_i converge to the desired circle domain.

Remarks 2.6. It was later proved that the Koebe Conjecture 2.3 holds for countably connected domains, that is, domains having countably many boundary components, see [HS93, Theorem 0.1]. Incidentally, the proof uses our main tool, the fixed-point index, in particular the Circle Index Lemma 3.1. Proofs of many other special cases of the KC 2.3 are known, and the same article [HS93] contains a list of references.

There are counterexamples to the uniqueness part of the statement of the KC 2.3 for uncountably connected domains. A brief discussion of this appears in the paragraph following the statement of Theorem 0.1 in the same article. The existence part in this case is still open.

The Circle Packing Theorem 2.2 is then proved as follows. Fix an embedding of G in the plane. For every vertex v pick a domain Ω_v bounded by a Jordan curve γ_v so that v is contained in the domain Ω_v, but no other vertex is, and so that the Ω_v are disjoint. Let Ω be the domain obtained by removing the closures of the Ω_v from $\hat{\mathbb{C}}$. Then by Koebe's Theorem 2.5, the domain Ω is biholomorphically equivalent to a circle domain, call it $\circledcirc_1$. Repeat this process, ensuring at each successive iteration that if $\langle u, v\rangle$ is an edge in G, then the boundaries of Ω_u and Ω_v have moved closer together. In this way we obtain a sequence of circle domains $\circledcirc_i$. One then shows that under the appropriate normalizations, the complements of the $\circledcirc_i$ converge to the desired circle packing.

Thurston has an elementary constructive existence proof, which involves adjusting the radii of the circles iteratively. He writes, "A change of any single radius affects most strongly the curvature at that vertex, so this process converges reasonably well," see [Thu80, Section 13.7]. Thurston's method is implemented in Stephenson's program `circlepack`, available online at `http://www.math.utk.edu/~kens/CirclePack/`, to draw circle packings having prescribed contact graphs.

2.2 Rigidity

Given the Circle Packing Theorem 2.2, we have as a natural follow-up question:

Question 2.7. *To what extent is the circle packing realizing a given graph unique?*

We say that $\mathcal{P}$ *realizes* G if G is the contact graph of the circle packing $\mathcal{P}$. Of course, Möbius and anti-Möbius transformations send circle packings to circle packings while preserving the contact graph, so the best we can hope for is uniqueness up to action by these maps. It is easy to construct examples of circle packings which have the same contact graph, but which are not Möbius equivalent. However, the following holds:

Koebe–Andreev–Thurston Theorem 2.8. *Let X be a triangulation of the 2-sphere $\mathbb{S}^2$. Then the circle packing realizing X is unique, up to action by Möbius and anti-Möbius transformations.*

A *triangulation* of a surface S is a collection of topological triangles with edge identifications so that the resulting object is homeomorphic to S. Thus in particular any triangulation of a surface without boundary is locally finite, and any triangulation of a compact surface without boundary is globally finite. To avoid pathological cases we require that

- two distinct faces of a triangulation X may only meet along a single edge, or at a single vertex, or not at all, and that
- there are no edge nor vertex identifications along the boundary of any one face of X.

We consider triangulations only up to their combinatorics. When we say that a circle packing *realizes* a triangulation of the 2-sphere $X = (V, E, F)$, we mean that it realizes the graph (V, E) which is the 1-skeleton of the triangulation.

We remark at this point that if G is a finite simple planar graph which is not the 1-skeleton of a triangulation of the 2-sphere, then there exist circle packings in $\hat{\mathbb{C}}$ realizing G which are not Möbius equivalent. Thus the uniqueness statement given in the KATT 2.8 is the best possible.

There is an easy proof of the KATT 2.8 based on a discrete version of the Discrete Complex Maximum Modulus Principle 2.9. We give this proof now to build intuition, and because we will refer back to the ideas of the proof in the future. This proof can be found in [Roh11, Section 2.4.3], where it is attributed to Oded Schramm. The earliest published version of this argument that we are aware of appears in [He99, Section 2].

We first state and prove the DCMMP 2.9. If D is a metric closed disk in $\mathbb{C}$ then let $R_{\mathbb{C}}(D)$ denote its Euclidean radius.

Discrete Maximum Modulus Principle 2.9. *Let D, D_i, be pairwise interiorwise disjoint metric closed disks in $\mathbb{C}$, with $i \in \mathbb{Z}/n\mathbb{Z}$ and $n \geq 3$. Suppose that the D_i form a tangent cycle around D. More precisely, suppose that D and D_i are tangent for all i, and that D_i and D_{i+1} are tangent for all i. Let $\tilde{D}, \tilde{D}_i$ also be pairwise interiorwise disjoint metric closed disks in $\mathbb{C}$ satisfying the analogous conditions. Then, unless $R_{\mathbb{C}}(D)/R_{\mathbb{C}}(\tilde{D}) = R_{\mathbb{C}}(D_i)/R_{\mathbb{C}}(\tilde{D}_i)$ for all i, there exists some i so that $R_{\mathbb{C}}(D)/R_{\mathbb{C}}(\tilde{D}) < R_{\mathbb{C}}(D_i)/R_{\mathbb{C}}(\tilde{D}_i)$.*

Proof. This is a consequence the following simple observation:

Observation 2.10. *Let z and z_i denote the Euclidean centers of D and D_i, respectively. Then the acute angle $\measuredangle z_i z z_{i+1}$ is monotone increasing in both $R_{\mathbb{C}}(D_i)$ and $R_{\mathbb{C}}(D_{i+1})$, and is monotone decreasing in $R_{\mathbb{C}}(D)$.*

Then suppose for contradiction that $R_{\mathbb{C}}(D)/R_{\mathbb{C}}(\tilde{D}) > R_{\mathbb{C}}(D_i)/R_{\mathbb{C}}(\tilde{D}_i)$ for some i, and simultaneously that there is no i so that $R_{\mathbb{C}}(D)/R_{\mathbb{C}}(\tilde{D}) < R_{\mathbb{C}}(D_i)/R_{\mathbb{C}}(\tilde{D}_i)$. Scale the collection of disks $\tilde{D}, \tilde{D}_i$ so that $R_{\mathbb{C}}(D) = R_{\mathbb{C}}(\tilde{D})$. Then for every i we have that $R_{\mathbb{C}}(D_i) \leq R_{\mathbb{C}}(\tilde{D}_i)$, with strict inequality for some i. Let $\tilde{z}$ and $\tilde{z}_i$ denote the Euclidean centers of $\tilde{D}$ and $\tilde{D}_i$, respectively. Then by Observation 2.10:

$$\sum_{i \in \mathbb{Z}/n\mathbb{Z}} \measuredangle z_i z z_{i+1} < \sum_{i \in \mathbb{Z}/n\mathbb{Z}} \measuredangle \tilde{z}_i \tilde{z} \tilde{z}_{i+1}$$

However, because the D_i form a tangent cycle around D, similarly the $\tilde{D}_i$ around $\tilde{D}$, both of these sums must be 2π, so we obtain our contradiction. □

Proof of the Koebe–Andreev–Thurston Theorem 2.8. Let $\mathcal{P} = \{D_v\}$ and $\tilde{\mathcal{P}} = \{\tilde{D}_v\}$ be two circle packings in $\hat{\mathbb{C}}$ realizing the same triangulation $X = (V, E, F)$ of the 2-sphere $\mathbb{S}^2$. The following is well-known, and not hard to check:

Fact 2.11. *Suppose that z_1, z_2, z_3 are distinct points on the Riemann sphere $\hat{\mathbb{C}}$. Then there is a unique triple of metric closed disks in $\hat{\mathbb{C}}$ having pairwise disjoint interiors, so that the three disks are pairwise tangent at the points z_i.*

Pick a triangle $\langle v_1, v_2, v_3 \rangle \in F$. Then we may assume without loss of generality by applying a Möbius transformation and possibly a reflection that $D_i = \tilde{D}_i$ for $i = 1, 2, 3$. Normalize further so that the D_i are metric closed disks in the plane, and so that all other disks of the two packings $\mathcal{P}$ and $\tilde{\mathcal{P}}$ are contained in the bounded region formed by the $D_i = \tilde{D}_i$ for $i = 1, 2, 3$. We will be done once we show that D_v and $\tilde{D}_v$ have the same radii for every vertex $v \in V$ of X.

Let v be a vertex of X where $R_{\mathbb{C}}(D_v)/R_{\mathbb{C}}(\tilde{D}_v)$ is maximized. Such a v exists because X has finitely many vertices. If this ratio equals 1 for every $v \in V$ then we are done, so suppose for contradiction, after interchanging the roles of $\mathcal{P}$ and $\tilde{\mathcal{P}}$ if necessary, that this ratio exceeds 1 for some $v \in V$. This implies $v \neq v_1, v_2, v_3$ because $R_{\mathbb{C}}(D_i)/R_{\mathbb{C}}(\tilde{D}_i) = 1$ for $i = 1, 2, 3$.

Then by the DCMMP 2.9 if w is any neighbor of v in the 1-skeleton of X, then $R_{\mathbb{C}}(D_w)/R_{\mathbb{C}}(\tilde{D}_w) = R_{\mathbb{C}}(D_v)/R_{\mathbb{C}}(\tilde{D}_v)$, so in particular the ratio we are considering is

maximized at w also. So, by stepping from vertex to vertex we get $R_{\mathbb{C}}(D_w)/R_{\mathbb{C}}(\tilde{D}_w) = R_{\mathbb{C}}(D_v)/R_{\mathbb{C}}(\tilde{D}_v) > 1$ for every $w \in V$, contradicting $D_1 = \tilde{D}_1$. □

Koebe gave a statement equivalent to the KATT 2.8 in [Koe36], which we discuss later in Section 2.11. The CPT 2.2 and the KATT 2.8 were rediscovered by Thurston in the late 1970s. He first brought these theorems to the attention of the mathematics community at large in his talk at the International Congress of Mathematicians, Helsinki, 1978, according to [Sac94, p. 135]. There he discussed his methods of proof based essentially on results of Andreev concerning finite-volume hyperbolic polyhedra, see [And70]. Andreev's results are too technical to state here, although we discuss a powerful generalization of his results due to Rivin in Section 2.10. The best source we are aware of for Thurston's work on this topic is his widely circulated lecture notes, [Thu80, Chapter 13]. There he notes that his methods can be generalized without too much trouble to give stronger results, discussed in Section 2.10. He also observes that the rigidity of the KATT 2.8 can alternatively be obtained as a consequence of the finite-volume Mostow–Prasad Rigidity Theorem 2.58; we describe how in Section 2.11. Today many proofs exist of existence and uniqueness of packings having contact graphs triangulating the sphere. For example, some proofs using variational principles appear in [CdV91, Brä92, Riv94].

2.3 Discrete Riemann mapping

This section is not directly related to the new mathematical results of this thesis, but is here for reasons of exposition and motivation. A central theme is that many notions from classical complex analysis have as their "discrete analogs" statements about circle packings. We saw an example of this already, in the proof of the KATT 2.8 via an application of a "discrete maximum principle." Perhaps the most important example of this general theme is that the KATT 2.8 can be interpreted as the "discrete analog of the classical Riemann Mapping Theorem 2.4." We first state a corollary to the CPT 2.2 and KATT 2.8:

Corollary 2.12. *Let $X = (V, E, F)$ be a triangulation of a topological closed disk. Then there exists a circle packing $\mathcal{P} = \{D_v\}$ realizing X in the closed unit disk $\bar{\mathbb{D}} \subset \mathbb{C}$ by Euclidean metric disks D_v, so that v is a boundary vertex of X if and only if D_v is internally tangent to $\partial\mathbb{D}$. Furthermore $\mathcal{P}$ is uniquely determined by X, up to action by Möbius and anti-Möbius transformations fixing $\mathbb{D}$ set-wise.*

For $\varepsilon > 0$ let $\mathcal{P}_{6,\varepsilon}$ denote the "hexagonal packing," also known as the "penny packing," the infinite constant-degree-6 packing of the plane by disks of constant

radius, in this case ε. See Figure 1.10a on p. 8. Here the *degree* of a disk is how many neighbors it has in the packing.

Pick $\Omega \subsetneq \mathbb{C}$ to be a simply connected open set. For $\varepsilon > 0$ consider the set $\mathcal{D}_\varepsilon$ of disks in $\mathcal{P}_{6,\varepsilon}$ which are completely contained in Ω. For ε sufficiently small the contact graph of $\mathcal{D}_\varepsilon$ is a triangulation X_ε of a topological closed disk. Then let $\mathcal{P}_\varepsilon$ be a circle packing in $\bar{\mathbb{D}}$ realizing X_ε, as per Corollary 2.12. We have a natural bijection between the disks of $\mathcal{D}_\varepsilon$, lying in Ω, and those of $\mathcal{P}_\varepsilon$, lying in $\mathbb{D}$. Define a function $\varphi_\varepsilon : \Omega \to \mathbb{D}$ "roughly" according to this bijection, in the natural way. Then:

Rodin–Sullivan Theorem 2.13. *Suitably normalized, the φ_ε converge to the Riemann mapping $\Omega \to \mathbb{D}$ as $\varepsilon \to 0$.*

This theorem was originally posed as a conjecture by Thurston, in his address at the International Symposium in Celebration of the Proof of the Bieberbach Conjecture, Purdue University, March 1985, according to [HS93, p. 371]. His conjecture is often credited with opening up the field of circle packing. The conjecture was proved by Rodin and Sullivan in [RS87]. The second half of Stephenson's book [Ste05] is devoted to the study of the connections between circle packing and classical complex analysis. In particular, he proves the Rodin–Sullivan Theorem 2.13 in Chapter 19. The Rodin–Sullivan Theorem 2.13 has been generalized considerably. For further details and references, see [Ste05, Chapter 20].

In a similar vein, the Measurable Riemann Mapping Theorem on existence and uniqueness of quasi-conformal maps has been proved via circle packing in [He90], and independently in [Wil06].

2.4 Discrete uniformization

The next natural question to ask is:

Question 2.14. *Can we find some existence and uniqueness theorems for circle packings in the complex plane, or in the hyperbolic plane?*

This question, too, has a complete answer:

Discrete Uniformization Theorem 2.15. *Let X be a triangulation of a topological open disk. Then there exists a locally finite circle packing $\mathcal{P}$ in $\mathbb{G}$ realizing X, where $\mathbb{G}$ is one of $\mathbb{C}$ and $\mathbb{H}^2$. Furthermore $\mathbb{G}$ is uniquely determined by X. Also $\mathcal{P}$ is uniquely determined by X up to action by conformal and anti-conformal automorphisms of $\mathbb{G}$, that is, by Euclidean similarities if $\mathbb{G} = \mathbb{C}$ and by hyperbolic isometries if $\mathbb{G} = \mathbb{H}^2$.*

This is the circle packing analog of the classical Uniformization Theorem, c.f. [For91, Theorem 27.9]:

Uniformization Theorem 2.16. *Let M be a simply connected one-dimensional differentiable complex manifold. Then M is conformally isomorphic to exactly one of the Riemann sphere $\hat{\mathbb{C}}$, the complex plane $\mathbb{C}$, and the hyperbolic plane $\mathbb{H}^2$.*

Traditionally, to parallel the UT 2.16 more closely, the DUT 2.15 is stated for X a triangulation of a simply connected topological surface, in which case the statement is exactly the combination of the CPT 2.2 and the KATT 2.8 with the statement of the DUT 2.15 we have given. We discuss the history of the DUT 2.15 at the end of this section.

Throughout we will consider the Poincaré unit disk model of the hyperbolic plane $\mathbb{H}^2$. We will implicitly refer to the following fact frequently:

Fact 2.17. *If $D \subset \mathbb{H}^2$ is a hyperbolic metric disk, then the image of D in $\mathbb{C}$ under the Poincaré unit disk embedding $\mathbb{H}^2 \hookrightarrow\!\!\!\rightarrow \mathbb{D} \hookrightarrow \mathbb{C}$ is a Euclidean metric disk in $\mathbb{C}$. Of course, the image of the hyperbolic center of $D \subset \mathbb{H}^2$ is not the Euclidean center of $D \subset \mathbb{C}$, unless both happen to be the origin $0 \in \mathbb{C}$.*

We now work out an example in detail. The *degree* or *valence* of a vertex in a graph, or in a triangulation, is the number of edges incident to it.

Example 2.18. Let X_6 be the 6-valent triangulation of the plane, meaning the triangulation so that every vertex has degree 6. It is easy to see that there is only one, considered up to its combinatorics. Then the circle packing realizing X_6 is the "penny packing" $\mathcal{P}_6$ in $\mathbb{C}$. In particular, the DUT 2.15 implies that there is no packing realizing X_6 in the hyperbolic plane.

Next, let X_7 be the 7-valent triangulation of an open disk. Again, there is only one. Then, there is a locally finite circle packing $\mathcal{P}_7$ by constant-radius hyperbolic disks realizing X_7 in the hyperbolic plane. To see why, we construct such a packing.

It suffices to show that there exists $r \in \mathbb{R}$ so that hyperbolic disks of hyperbolic radius r may be arranged to perfectly surround another hyperbolic disk of radius r. We show that for r too small, the seven satellite disks wrap too far around the central disk, and for r too big, it is impossible for seven satellite disks to wrap all the way around the central disk. Then by continuity there is an intermediate r which works. Hyperbolic geometry "looks Euclidean" on a small scale, in the sense that small nearby disks of the same hyperbolic radius have similar Euclidean radii under the unit disk embedding $\mathbb{H}^2 \hookrightarrow\!\!\!\rightarrow \mathbb{D}$. Thus it is clear that if r is sufficiently small,

then our seven satellite disks will wrap too far around the central disk, because six Euclidean disks of constant radius wrap perfectly around a central disk having the same Euclidean radius. Next, suppose we take our central disk D to be centered at the origin $0 \in \mathbb{D}$. If r is sufficiently big, then no seven Euclidean disks tangent to D, but contained in the unit disk $\mathbb{D}$, can wrap around D.

Note that $\mathcal{P}_7$ is also a circle packing in $\mathbb{C}$ via the Poincaré embedding $\mathbb{H}^2 \hookrightarrow \mathbb{D}$, but $\mathcal{P}_7$ is not locally finite in $\mathbb{C}$. It turns out that the converse situation does not occur: if $\mathcal{P}$ is a locally finite circle packing in $\mathbb{C}$ realizing the triangulation X of an open disk, then there is no circle packing in $\mathbb{H}^2$ realizing X, locally finite or not.

We now sketch the proof of the existence part of the DUT 2.15. This construction is standard, and its details can be found in [Ste05, Chapter 8] or in [HS95, Theorem 7.1]. We will use Corollary 2.12. Fix X as in the statement of the DUT 2.15. Pick a vertex v_0 of X, and let $X_1, X_2, \ldots$ be an exhaustion of X by triangulations of a topological closed disk, all containing v_0, so that $\cup_{i=1}^{\infty} X_i = X$. For each X_i let $\mathcal{P}_i = \{D_v^i\}$ be a circle packing coming from Corollary 2.12, normalized so that $D_{v_0}^i \in \mathcal{P}_i$ is centered at the origin $0 \in \mathbb{D} \cong \mathbb{H}^2$ for all i. One shows first that the hyperbolic radii of the disks of the packings $\mathcal{P}_i$ is monotone decreasing under this procedure. This follows from the following lemma, which is the circle packing analog of the classical Schwarz–Pick Lemma:

Discrete Schwarz–Pick Lemma 2.19. *Let Y be a triangulation of a topological closed disk, and let $\mathcal{Q}$ be the packing realizing Y as in Corollary 2.12. Let $\tilde{\mathcal{Q}}$ be another circle packing realizing Y, so that every disk of $\tilde{\mathcal{Q}}$ is contained in the unit disk $\bar{\mathbb{D}}$. Then the (possibly infinite) hyperbolic radius of the disk $E_v \in \mathcal{Q}$ for the vertex v is at least the hyperbolic radius of $\tilde{E}_v \in \tilde{\mathcal{Q}}$. Furthermore, equality of these radii at even a single interior vertex, or equality at every boundary vertex, implies equality at every vertex.*

This was originally proved by a Perron method in [BS91b]. However, it can be proved easily by an application of the Discrete Complex Maximum Modulus Principle 2.9, as a corollary of the following observations:

Observation 2.20. *Let D_0 be a metric closed disk contained in $\mathbb{D} \cong \mathbb{H}^2$ centered at the origin, and let D be a metric closed disk with variable radius which is tangent to D_0. Then $R_{\mathbb{H}^2}(D)$, if finite, is monotone increasing in $R_{\mathbb{C}}(D)$. Furthermore for real $\delta > 1$, we have that $R_{\mathbb{H}^2}(\delta D)/R_{\mathbb{H}^2}(D) > R_{\mathbb{C}}(\delta D)/R_{\mathbb{C}}(D) = \delta$.*

Here δD is dilation of D about the origin, so $\delta D = \{\delta z : z \in D\}$. This argument is adapted from [He99, Lemma 2.2].

Thus by our normalization one of two cases occurs:

1. The Euclidean radius of $D^i_{v_0}$ converges to some positive real number.

2. The Euclidean radius of $D^i_{v_0}$ converges to 0.

In fact, we find that either case (1) occurs for all $v \in V$, or case (2) occurs for all $v \in V$. This follows from the following lemma, c.f. [RS87, p. 352; Ste05, Lemma 8.2]:

Discrete Ring Lemma 2.21 (Rodin, Sullivan)**.** *Suppose that $n \geq 3$ interiorwise disjoint closed disks D_i, encircle the unit disk $\mathbb{D}$, so that all are externally tangent to $\mathbb{D}$. Then there is a constant depending only on n that lower bounds the Euclidean radii of the encircling disks D_i.*

Attempting to draw a counterexample is generally enough to convince oneself of this lemma.

So, suppose that all of the Euclidean radii of the D^i_v converge to positive real numbers. Then the subsequential limiting object of the $\mathcal{P}_i$ is a circle packing $\mathcal{P}$ in the unit disk $\mathbb{D} \cong \mathbb{H}^2 =: \mathbb{G}$ realizing X. Next, suppose that the Euclidean radii of the D^i_v converge to 0. In this case one applies an additional normalization to all of the $\mathcal{P}_i$, scaling each so that the Euclidean radius of $D^i_{v_0}$ is equal to 1. Then the subsequential limiting object of the $\mathcal{P}_i$ is a circle packing $\mathcal{P}$ in $\mathbb{C} =: \mathbb{G}$ realizing X.

In either case, as the final step in establishing existence, one needs to show that $\mathcal{P}$ is locally finite in $\mathbb{G}$. This turns out to be non-trivial. This last part of the proof is worked out using totally different methods in our two main references. In [Ste05, Section 8.2], Stephenson uses a topological argument via our main tool, fixed-point index, and the Discrete Schwarz–Pick Lemma 2.19 for circle packings. We give further discussion on the Discrete Schwarz–Pick Lemma 2.19 shortly. His proof is inspired by similar arguments carried out in a much more general setting by He and Schramm in [HS93, Corollary 0.5], where a proof of the full statement of the DUT 2.15 is given via arguments related to the Koebe Conjecture 2.3. Stephenson attributes to this paper the main ideas of the proof he gives. On the other hand, in [HS95, Theorem 7.1], Z.-X. He and Schramm complete this existence proof by vertex extremal length arguments. Their existence statement is more general, in that it allows packing in arbitrary simply connected planar domains, not necessarily $\mathbb{D}$ or $\mathbb{C}$.

We next discuss the rigidity portion of the DUT 2.15. Rigidity of locally finite circle packings in $\mathbb{C}$ in the bounded-valence case is easily proved by a slight modification of an argument by Rodin and Sullivan in [RS87, Appendix 1]. Their proof

uses quasi-conformal mapping theory, and has no hope of being generalized to the unbounded-valence case. Schramm gave the first proof of the rigidity part of the DUT 2.15 in the unbounded valence case in [Sch91, Rigidity Theorems 1.1, 5.1]. His proofs are elementary, employing essentially only plane topology arguments and manipulations with the Möbius group. A different proof is given with a similar main idea in the article already mentioned [HS93] by He and Schramm. Similar rigidity theorems in more general settings have been proved using vertex extremal length, as in [He99]. We will discuss this more in Section 2.10.

We give a proof similar to Schramm's of the rigidity part of the DUT 2.15 when $\mathbb{G} = \tilde{\mathbb{G}} = \mathbb{C}$, using fixed-point index in Section 3.5. The proof we present comes from [Ste05, Section 8.3.3], and is adapted from arguments coming from [HS93]. In the same section, we adapt this proof to show that there cannot be two circle packings, one in $\mathbb{C}$ and the other in $\mathbb{H}^2$, both locally finite in their respective spaces, realizing the same triangulation of a topological open disk.

Before moving on, we outline the proof of rigidity portion of the DUT 2.15 when $\mathbb{G} = \tilde{\mathbb{G}} = \mathbb{H}^2$. This proof is adapted from the proof of [He99, Rigidity Theorem 1.2], and can also be found in [Ste05, Section 8.3.2]. First, if D is a disk which meets $\mathbb{D} \cong \mathbb{H}^2$ but is not contained in $\mathbb{D}$, then we formally write $R_{\mathbb{H}^2}(D) = \infty$, and adopt the convention that $r < \infty$ for any real number r. We need the following straightforward generalization of the Discrete Schwarz–Pick Lemma 2.19, which is also proved via the Discrete Maximum Modulus Principle 2.9:

Lemma 2.22. *Let Y be a triangulation of a topological closed disk, and let $\mathcal{Q} = \{E_v\}$ and $\tilde{\mathcal{Q}} = \{\tilde{E}_v\}$ be circle packings realizing Y. Suppose that all of the disks of $\mathcal{Q}$ meet $\mathbb{D}$, and all of the disks of $\tilde{\mathcal{Q}}$ are contained in $\mathbb{D}$. Then, if the inequality $R_{\mathbb{H}^2}(\tilde{E}_v) \leq R_{\mathbb{H}^2}(E_v)$ holds for all boundary vertices v, then it holds for all interior vertices v.*

For the details of the proof of Lemma 2.22, see [He99, Section 2].

Now, let $\mathcal{P}$ and $\tilde{\mathcal{P}}$ be circle packings realizing a shared triangulation X of a topological open disk, so that both $\mathcal{P}$ and $\tilde{\mathcal{P}}$ are locally finite in $\mathbb{H}^2 \cong \mathbb{D}$. Suppose for contradiction that there is a pair of disks $D_w \in \mathcal{P}$ and $\tilde{D}_w \in \tilde{\mathcal{P}}$ with unequal hyperbolic radii. Normalize so that both are centered at the origin $0 \in \mathbb{D} \subset \mathbb{C}$, and suppose without loss of generality that $R_{\mathbb{H}^2}(D_w) < R_{\mathbb{H}^2}(\tilde{D}_w)$. Because both are centered at the origin this is equivalent to $R_{\mathbb{C}}(D_w) < R_{\mathbb{C}}(\tilde{D}_w)$. Then let $\delta > 1$ so that $\delta R_{\mathbb{C}}(D_w) < R_{\mathbb{C}}(\tilde{D}_w)$. Let $\mathcal{P}_0$ be the set of disks of $\delta\mathcal{P} = \{\delta D : D \in \mathcal{P}\}$ which meet $\bar{\mathbb{D}}$. Then because $\mathcal{P}$ is locally finite in $\mathbb{D}$, the packing $\mathcal{P}_0$ is finite. Furthermore, by adding finitely many disks of $\delta\mathcal{P}$ if necessary, we may suppose also that the contact

graph of $\mathcal{P}_0$ is the 1-skeleton of a triangulation X_0 of a topological closed disk. Write D_v^0 to distinguish the disk for v in $\mathcal{P}_0$ from that in $\mathcal{P}$. Let $\tilde{\mathcal{P}}_0$ be the set of disks of $\tilde{\mathcal{P}}$ whose vertices belong to X_0. Then if v is a boundary vertex of X_0, we have that $R_{\mathbb{H}^2}(D_v^0)$ is infinite. On the other hand, because the disks of $\tilde{\mathcal{P}}_0$ lie inside of $\mathbb{D}$ we have that $R_{\mathbb{H}^2}(\tilde{D}_v)$ is finite. Thus by Lemma 2.22 we have that $R_{\mathbb{H}^2}(\tilde{D}_v) < R_{\mathbb{H}^2}(D_v^0)$ for all $v \in X_0$, thus $R_{\mathbb{C}}(\tilde{D}_w) < R_{\mathbb{C}}(D_w^0) = \delta R_{\mathbb{C}}(D_w)$, a contradiction.

2.5 Generalizing to other surfaces

Question 2.23. *Suppose that we start with a triangulation X of a topological surface S other than the sphere and the open disk. Is there then a circle packing, appropriately defined, realizing X, and is it unique in some nice sense?*

At this point there is ambiguity surrounding the phrase *circle packing*. In particular, we would like to have a metric structure on S so that we can talk about metric closed disks, as Question 2.23 is trivial if we consider packings by topological disks. Fortunately, and somewhat surprisingly, the combinatorics of X determine our metric structure for us, as per the following theorem:

Theorem 2.24. *Given a triangulation X of an oriented topological surface S, there is a complete constant curvature Riemannian metric d on S and a locally finite circle packing $\mathcal{P}$ in (S, d) so that $\mathcal{P}$ realizes X. Furthermore d is uniquely determined by X up to conformal and anti-conformal isomorphisms, and $\mathcal{P}$ is then uniquely determined by X and d up to conformal and anti-conformal automorphisms of (S, d).*

We discuss the history of this theorem at the end of the section. By *topological surface* we mean a connected Hausdorff 2-dimensional real manifold. Thus S may be open or closed. We do not consider S with boundary. As usual, a metric is *complete* if every Cauchy sequence converges.

By the Uniformization Theorem 2.16, the special case where S is simply connected is exactly the DUT 2.15. Next, recall the following corollary to the Uniformization Theorem:

Theorem 2.25. *Let R be a Riemann surface homeomorphic to S. Then S admits a complete constant curvature Riemannian metric d, uniquely determined by R up to conformal isomorphisms of (S, d), so that (S, d) is conformally isomorphic to R.*

Recall that a *Riemann surface* is a connected Hausdorff 1-dimensional differentiable complex manifold. Thus Theorem 2.24 is a corollary of the DUT 2.15, by the following

argument: let X be a triangulation of the non-simply-connected surface S. Then X lifts to a periodic triangulation X_{U} of S_{U} the universal cover of S. We then obtain a circle packing $\mathcal{P}_{\mathrm{U}}$ realizing X_{U} in $(S_{\mathrm{U}}, d_{\mathrm{U}})$. It is not hard to show that $\mathcal{P}_{\mathrm{U}}$ is periodic because X_{U} is, so we may obtain our desired circle packing $\mathcal{P}$ realizing X in the quotient space (X, d). The uniqueness part follows from uniqueness in the simply connected case. We remark that the case of unoriented S may be taken care of using a similar argument, by taking the oriented double cover of S.

The first half of Stephenson's book [Ste05] consists of a complete and self-contained proof of Theorem 2.24, from the CPT 2.2 and KATT 2.8 to the DUT 2.15. At many crucial points he relies on the methods of Schramm and He. In [BS90] a statement and self-contained proof is given of Theorem 2.24 with the significant additional hypothesis that the vertex degree of X is uniformly bounded above; they use this hypothesis strongly in their proof of the (bounded-valence) DUT 2.15, where their argument appeals to the theory of quasi-conformal mapping. Another statement similar to the DUT 2.15 is given and proved in [HS93, Corollary 0.5].

2.6 What is a thin disk configuration?

The new results contained in this thesis stem from the following question:

Question 2.26. *If we consider collections of disks in which we allow overlaps, to what extent can we recover our existence and uniqueness statements?*

We are not the first to study collections of disks with overlaps. A discussion of related results is given in Section 2.10.

More precisely, we use the term *disk configuration* to mean a collection of closed disks, none of which is contained in any other, but with no other conditions. In this thesis we focus on a particular class of disk configurations, namely the *thin* ones:

Definition 2.27. A disk configuration is called *thin* if no two of its disks meet in the interior of a third.

For two examples of thin disk configurations see Figure 1.10 on p. 8.

2.6.1 Recognizing thin disk configurations

Let $\mathcal{C} = \{D_v\}_{v \in V}$ be a disk configuration, and let $G = (V, E)$ be its contact graph, defined as before. Let $\Theta : E \to [0, \pi)$ be a function. Then we say that $\mathcal{C}$ *realizes* (G, Θ) if for every edge $\langle u, v \rangle$ of G, equivalently for every pair $D_u, D_v \in \mathcal{C}$ which meet,

we have that $\measuredangle(D_u, D_v) = \Theta(\langle u, v\rangle)$, c.f. Figure 2.2. We use the term *incidence data* to refer to such a pair (G, Θ), consisting of the graph $G = (V, E)$ and the function $\Theta : E \to [0, \pi)$.

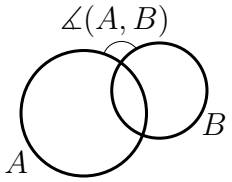

Figure 2.2: **The definition of** $\measuredangle(A, B)$ the *external intersection angle* or *overlap angle* between two closed disks A and B.

Our first observation is that it is easy to check whether a disk configuration is thin given only its incidence information:

Proposition 2.28. *Let $\mathcal{C}$ be a disk configuration in $\hat{\mathbb{C}}$ realizing the incidence data (G, Θ). Then $\mathcal{C}$ is thin if and only if the following condition holds:*

- *Suppose that three edges e_1, e_2, e_3 of G form a 3-cycle in G. Then we require that $\Theta(e_1) + \Theta(e_2) + \Theta(e_3) \leq \pi$.*

The main idea of this proposition is essentially contained in Figure 2.3. We leave checking the details to the reader. As a consequence, if two disk configurations $\mathcal{C}$ and $\tilde{\mathcal{C}}$ in $\hat{\mathbb{C}}$ realize the same incidence data (G, Θ), then $\mathcal{C}$ is thin if and only $\tilde{\mathcal{C}}$ is.

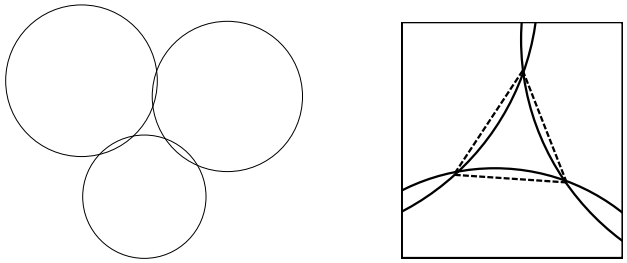

Figure 2.3: **A triple of disks meeting thinly, and the sum of their overlap angles.** If we zoom in on the "hole" bounded by the three disks, we see that the sum of their overlap angles is at most the sum of the angles in the dashed triangle.

The condition of thinness limits the combinatorial situations we need to consider, and we use it strongly in our main technical result, the Index Theorem 4. Our main focus in this thesis is to prove via elementary means a rigidity statement for configurations of disks where we allow overlaps. For completeness, we include a discussion on the question of existence in Section 2.9.

Before stating our main results in Section 2.7, we require a discussion on the types of contacts that can occur in a thin disk configuration, contained in Section 2.6.2.

2.6.2 Mostly-triangulations

We have seen that a circle packing on a surface S gives a natural embedding of its contact graph into S. However, the contact graphs of arbitrary disk configurations are not characterized so easily. For example, it is easy to see that there are disk configurations in the plane $\mathbb{C}$ whose contact graphs are not planar. In particular, consider the configuration having at least 5 disks all of which meet the origin, none of which is contained in any other. In this section, we ask:

Question 2.29. *If $\mathcal{C}$ is a thin disk configuration, what combinatorial conditions characterize its contact graph?*

This question turns out to have a relatively simple answer. We first note:

Observation 2.30. *If $\mathcal{C}$ is a thin disk configuration, then at most four disks of $\mathcal{C}$ can meet at any one point, and the only way that four disks can meet at a point is as shown in Figure 2.4.*

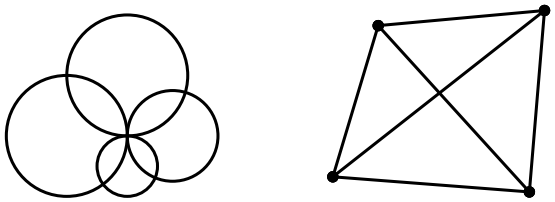

Figure 2.4: **Four closed disks meeting at a point, and the associated contact graph.**

We remark at this point that this control over the simultaneous intersection of disks of our configurations guaranteed by the thinness condition is an essential ingredient in our rigidity proofs. Anyway, with Observation 2.30 in mind, we make the following definition:

Definition 2.31. Let $G = (V, E)$ be a graph. Suppose that there is a collection $\mathcal{K}$ of subgraphs K_i^4 of G so that:

- each K_i^4 is a complete graph on 4 vertices, meaning that every two vertices in K_i^4 have an edge between them,
- no pair K_i^4 and K_j^4 with $i \neq j$ share more than 2 vertices, and
- if for each i, we delete from G an edge of K_i^4 which does not belong to any other K_j^4, then the resulting graph G_0 is embeddable in the topological surface S.

In that case we say that G is *mostly-embeddable* in S. We refer to the graph G_0 obtained after our edge deletions as an *S-reduction* of G. Finally, we say that $\psi : G \to S$ is a *mostly-embedding* of G in S if there is an S-reduction G_0 of G so that:

- the restriction of ψ to G_0 is an embedding in S, and
- for $e_1, e_2 \in E$, the embedded edges $\psi(e_1)$ and $\psi(e_2)$ cross if and only if one of e_1 and e_2 was deleted from G to obtain the S-reduction G_0.

Then we say that $\psi : G \to S$ is a *mostly-embedding for this S-reduction G_0*.

Before moving on, we make two remarks on this highly technical definition. First, every graph G that is embeddable in S is mostly-embeddable in S. Second, there may in general be many distinct S-reductions of G, which admit different mostly-embeddings.

By Observation 2.30, we get:

Observation 2.32. *If $\mathcal{C}$ is a thin disk configuration living in (S, d), where S is a surface and d is a Riemannian metric on S, then the contact graph G of $\mathcal{C}$ is mostly-embeddable in S.*

In fact, if we place a vertex for every disk D_v of $\mathcal{C}$ at the metric center of D_v, and connect neighboring disks D_u and D_v by the geodesic passing through their overlap region $D_u \cap D_v$, then the resulting representation of G is a mostly-embedding. Thus mostly-embeddability of G in some surface S is a necessary condition for G to appear as the contact graph of a thin disk configuration. Conversely:

Fact 2.33. *If G is mostly-embeddable in S then there exists some thin disk configuration $\mathcal{C}$ having contact graph G in (S, d) for some Riemannian metric d, although* a priori *we are not allowed to specify the overlap angles of the disks of $\mathcal{C}$.*

This is not obvious, but nevertheless serves as good motivation for Definition 2.31. It follows from Theorem 2.55.

Next:

Observation 2.34. *Suppose that G_0 is the 1-skeleton of a triangulation of a topological surface S. Then it is not hard to see that G_0 admits precisely one embedding in S, up to action by homeomorphisms $S \to S$.*

Let S be a surface, and suppose that G is a graph which admits an S-reduction G_0. Let $\psi : G \to S$ be a mostly-embedding for G_0. Suppose that for any S-reduction G_0' of G with associated mostly-embedding ψ', the following hold:

- We have that ψ restricts on G_0' to an embedding, similarly ψ' on G_0, and furthermore that
- the restrictions of ψ and of ψ' to $G_0 \cap G_0'$ differ by post-composition by a homeomorphism $S \to S$.

Then we say that G admits an *essentially unique* mostly-embedding $\psi : G \to S$. In that case, if $e_1, e_2 \in E$ are so that $\psi(e_1)$ and $\psi(e_2)$ cross, then one of e_1 and e_2 is deleted in any S-reduction of G. Then the following definition is well-stated as a consequence of Observation 2.34:

Definition 2.35. Suppose that the graph G admits an S-reduction G_0 which is the 1-skeleton of a triangulation of S. Then any other S-reduction of G is also the 1-skeleton of a triangulation of S, and G admits an essentially unique mostly-embedding $\psi : G \to S$. In this case we say that G is a *mostly-triangluation* of S.

2.7 Statements of the main results

Our main results are the following three theorems:

Main Rigidity Theorem 1. *Let $\mathcal{C}$ be a thin disk configuration in $\hat{\mathbb{C}}$ realizing incidence data (G, Θ). Suppose that G mostly-triangulates the 2-sphere. Suppose $\tilde{\mathcal{C}}$ is another disk configuration in $\hat{\mathbb{C}}$ realizing (G, Θ). Then $\mathcal{C}$ and $\tilde{\mathcal{C}}$ differ by a Möbius or anti-Möbius transformation.*

Main Rigidity Theorem 2. *Let $\mathcal{C}$ be a locally finite thin disk configuration in $\mathbb{C}$ realizing incidence data (G, Θ). Suppose that G mostly-triangulates a topological open disk. Suppose $\tilde{\mathcal{C}}$ is another locally finite disk configuration in $\mathbb{C}$ realizing (G, Θ). Then $\mathcal{C}$ and $\tilde{\mathcal{C}}$ differ by a Euclidean similarity.*

Main Uniformization Theorem 3. *There do not exist two thin disk configurations $\mathcal{C}$ locally finite in $\mathbb{C}$ and $\tilde{\mathcal{C}}$ locally finite in $\mathbb{H}^2$ realizing the same incidence data.*

Theorems 1, 2, and 3 are proved roughly the same way, using our technical result the Main Index Theorem 4. The proofs are in Section 4.3.

We remark that our Main Rigidity Theorem 1 may be obtained from Rivin's Theorem 2.45 via a construction described in Section 2.9.3, but our other two main theorems cannot. In general, the Riemann sphere $\hat{\mathbb{C}}$ is the setting in which rigidity statements of this type are easiest to prove, and the complex plane is the setting where doing so is most difficult.

We also state the following analogous rigidity result in the hyperbolic plane:

Theorem 2.36. *Let $\mathcal{C}$ be a locally finite thin disk configuration in $\mathbb{H}^2$ realizing incidence data (G, Θ). Suppose that G mostly-triangulates a topological open disk. Suppose $\tilde{\mathcal{C}}$ is another locally finite disk configuration in $\mathbb{H}^2$ realizing (G, Θ). Then $\mathcal{C}$ and $\tilde{\mathcal{C}}$ differ by a hyperbolic isometry.*

This can be proved by a modification of the proof in Section 2.4, where we proved rigidity of circle packings locally finite in $\mathbb{H}^2$ having contact graphs triangulating a topological open disk. In particular we would require a generalization of the Discrete Schwarz–Pick Lemma 2.19. It is straightforward to adapt the proof of the generalization of the same lemma given in [He99, Lemma 2.2] to our setting. We omit the details.

In a way, the main contribution of this thesis is that our proof via our Index Theorem 4 confirms that the elementary arguments for rigidity of circle packings in the plane given in [Sch91], and implicitly in [HS93], can be generalized to case where we allow some overlaps among our disks. It was previously thought that those arguments depended too crucially on pairwise interiorwise disjointness of the disks for there to be a hope of generalizing them in this direction, see comments in [He99, p. 3].

Our Main Theorems 1, 2, 3, together with Theorem 2.36, imply a uniqueness statement for thin disk configurations on other surfaces along the lines of the uniqueness portion of the Discrete Uniformization Theorem 2.15, and via an identical proof.

2.8 Conjectured generalizations of our main results

Most simply, we make the following conjecture:

Conjecture 2.37. *Our Main Theorems 1, 2, and 3 and Theorem 2.36 continue to hold with the thinness condition omitted.*

Our Main Index Theorem 4 has thinness as a hypothesis, and we will conjecture that we may eliminate the thinness condition there as well. It seems that then our proofs should generalize to establish the first three theorems listed in Conjecture 2.37. Then it seems implausible that the generalized version of Theorem 2.36 could fail if the generalizations of Theorems 1, 2, and 3 hold.

Even more strongly, we make two conjectures which together would subsume all other currently known rigidity and uniformization statements on disk configurations. First:

Conjecture 2.38. *Suppose that $\mathcal{C}$ and $\tilde{\mathcal{C}}$ are disk configurations, locally finite in $\mathbb{G}$ and $\tilde{\mathbb{G}}$ respectively, where each of $\mathbb{G}$ and $\tilde{\mathbb{G}}$ is equal to one of $\mathbb{C}$ and $\mathbb{H}^2$, with the* a priori *possibility that $\mathbb{G} \neq \tilde{\mathbb{G}}$. Suppose that $\mathcal{C}$ and $\tilde{\mathcal{C}}$ share a contact graph $G = (V, E)$. Suppose further that there is a planar subgraph of G which is a polytopal decomposition of a topological open disk. Then $\mathbb{G} = \tilde{\mathbb{G}}$.*

A *polytopal decomposition* of a topological surface is a decomposition of the surface into cells, each of which has finitely many sides, so that each vertex is incident to finitely many edges, but with no *a priori* uniform bounds on the degrees of either the faces or the vertices. We insist on the same regularity conditions as those we imposed on triangulations, c.f. p. 20. We remark that a polytopal decomposition of the 2-sphere $\mathbb{S}^2$ is exactly a combinatorial polyhedron. A natural approach to try to prove Conjecture 2.38 is via vertex extremal length arguments, along the lines of [He99, Uniformization Theorem 1.3] and [HS95]. Second:

Conjecture 2.39. *Suppose that $\mathcal{C}$ and $\tilde{\mathcal{C}}$ are disk configurations, both locally finite in $\mathbb{G}$, where $\mathbb{G}$ is equal to one of $\mathbb{C}$ and $\mathbb{H}^2$. Suppose that $\mathcal{C}$ and $\tilde{\mathcal{C}}$ realize the same incidence data (G, Θ). Suppose further that some maximal planar subgraph of G is the 1-skeleton of a triangulation of a topological open disk. Then $\mathcal{P}$ and $\tilde{\mathcal{P}}$ differ by a Euclidean similarity if $\mathbb{G} = \mathbb{C}$ or by a hyperbolic isometry if $\mathbb{G} = \mathbb{H}^2$.*

Conjectures 2.38 and 2.39 could follow from suitable generalizations of our technical theorems. We also make the natural conjecture analogous to Conjecture 2.39 for disk configurations on the Riemann sphere.

Finally, we conjecture that Conjecture 2.39 is the best possible uniqueness statement of its type, in the following precise sense:

Conjecture 2.40. *Let $\mathcal{C}$ be a disk configuration which is locally finite in $\mathbb{G}$, where $\mathbb{G}$ is one of $\hat{\mathbb{C}}$, $\mathbb{C}$, or $\mathbb{H}^2$. Suppose that $\mathcal{C}$ realizes incidence data (G, Θ). Suppose that no maximal planar subgraph of G is the 1-skeleton of a triangulation of a topological open disk. Then there are disk configurations in $\mathbb{G}$ realizing (G, Θ) which are not images of $\mathcal{C}$ under any conformal or anti-conformal automorphism of $\mathbb{G}$.*

2.9 The existence question for disk configurations

The following is a natural question, given our previous discussion:

Question 2.41. *Let G be a graph, and let $\Theta : E \to [0, \pi)$. Can we give necessary and sufficient conditions on G and Θ under which there exists a disk configuration realizing the incidence data (G, Θ)?*

No complete answer to this question is currently known. Considering disk configurations on multiply-connected surfaces complicates matters, so for the sake of simplicity, for the remainder of our discussion of Question 2.41 we restrict our attention to disk configurations in $\hat{\mathbb{C}}$. We will see an example of how things change on multiply connected surfaces in Remark 2.42.

It is easy to see that certain conditions are necessary for a positive answer to Question 2.41. For example:

($\measuredangle 1$) Suppose that v_i, $i \in \mathbb{Z}/n\mathbb{Z}$ are vertices forming an n-cycle in G, with $n \geq 4$, so that v_i and v_{i+1} share an edge in G, but v_i and v_j do not for $i \neq j \pm 1$. Then, we require that $\sum_{i=1}^{n} \Theta(\langle v_i, v_{i+1} \rangle)$ is strictly less than $(n-2)\pi$ the sum of the interior angles of a Euclidean n-gon.

To see why this condition is necessary, see for example Figures 1.16 and 1.17 on p. 12. The essential point is that certain contacts and overlap angles demanded by G and Θ can force additional contacts. So, for a positive answer to Question 2.41, these forced additional contacts, and perhaps some special angle conditions on them, should be reflected in (G, Θ). It seems that in general, finding clean statements of these conditions that force extra contacts, along the lines of ($\measuredangle 1$), is not easy. Our setting of thin disk configurations is somewhat more tractable, and is discussed next in Section 2.9.1.

Remark 2.42. A disk configuration on a multiply-connected surface S may realize a pair (G, Θ) which fails to satisfy condition ($\measuredangle 1$). In this case it turns out that we need to consider, in condition ($\measuredangle 1$), only those n-cycles whose edges' images, under the natural geodesic embedding, form a contractible closed loop in S. Similar technical modifications of the conditions described in Section 2.9.1 would be necessary if we wished to accommodate disk configurations in multiply-connected surfaces. For simplicity we do not mention this issue from now on.

2.9.1 Existence of thin disk configurations

Our main result is a rigidity theorem, and it would be nice to have a characterization, or at least an existence statement, for the collection of objects for which we are proving rigidity. Thin disk configurations are easy to draw directly, so clearly the collection of objects we are considering is non-trivial, and in fact quite large. We ask the following question, which we will see is a special case of Question 2.41:

Question 2.43. *Let G be a graph, and let $\Theta : E \to [0, \pi)$. Can we give necessary and sufficient conditions on G and Θ under which there exists a thin disk configuration realizing (G, Θ)?*

We remind the reader that for the remainder of the section, we consider only disk configurations in $\hat{\mathbb{C}}$. We first remark that a complete answer to the more general Question 2.41 would give a complete answer to Question 2.43 as well, via Proposition 2.28, which guarantees that a disk configuration realizing the incidence data (G, Θ) is thin if and only if the following condition on (G, Θ) holds:

($\measuredangle 2$) Suppose that three edges e_1, e_2, e_3 of G form a 3-cycle in G. Then we require that $\Theta(e_1) + \Theta(e_2) + \Theta(e_3) \leq \pi$.

We do not have a complete answer for Question 2.43. We describe several necessary conditions on (G, Θ), for a positive answer to the question in Section 2.9.2. Afterward, we state Theorem 2.44, which gives sufficient conditions for a positive answer to the same question. It is easily seen that these sufficient conditions are not necessary for a positive answer. Finally, we give a heuristic argument for why a complete answer to Question 2.43 may be hard to present, via a worked example in Section 2.9.6.

2.9.2 Some necessary conditions on (G, Θ) and an existence statement

The following conditions are easily seen to be necessary for the existence of a thin disk configuration in $\hat{\mathbb{C}}$ realizing (G, Θ):

($\measuredangle 3$) Suppose that G has an essentially unique mostly-embedding $\psi : G \to \mathbb{S}^2$, and that the edges e_1, e_2, e_3 of G form a 3-cycle in (V, E) that separates the graph (V, E), meaning that the graph obtained from G by removing the closed edges e_1, e_2, e_3 is disconnected. Then we require that $\sum_{i=1}^{3} \Theta(e_i) < \pi$.

This can be thought of as the analog of ($\measuredangle 2$) for $n = 3$. To see why, see Figure 2.5. Next:

($\measuredangle 4$) Suppose that the edges e_1, e_2, e_3, e_4 of $G = (V, E)$ form a 4-cycle in G in that order, and that $\sum_{i=1}^{4} \Theta(e_i) = 2\pi$. Then we require that the vertices which are the endpoints of the e_i form a complete subgraph in G. In particular, this together with the thinness condition on (G, Θ) implies two things: (1) for e equal to either of the two other edges of this subgraph besides e_1, e_2, e_3, e_4, we get that $\Theta(e) = 0$, and (2) we get $\Theta(e_1) = \Theta(e_3) = \pi - \Theta(e_2) = \pi - \Theta(e_4)$.

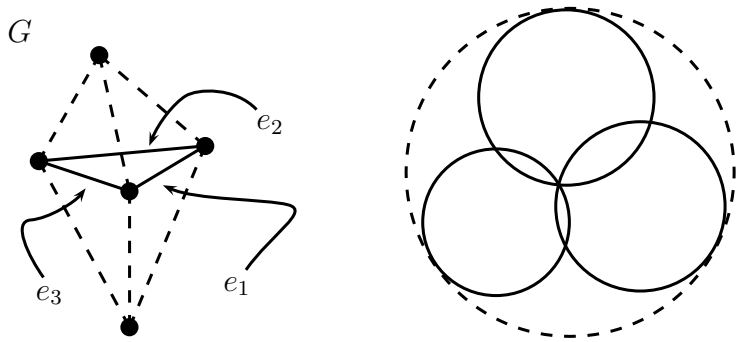

Figure 2.5: **Justification for angle condition (∠3).** Consider the graph $G = (V, E)$ as shown. It has an essentially unique embedding in $\hat{\mathbb{S}}^2$. Suppose we label the solid edges e_1, e_2, e_3 as shown. Define $\Theta : E \to [0, \pi)$ so that $\Theta(e_1)+\Theta(e_2)+\Theta(e_3) = \pi$, but $\Theta(e) = 0$ for all other edges e of G. The black disks are meant to represent the endpoints of the edges e_1, e_2, e_3: they will meet as shown. Then to realize (G, Θ), we need to add two disjoint disks, each of which is tangent to every solid disk. We cannot do so, because there is no "hole" in the middle of the solid disks, as there would be if we had instead $\Theta(e_1) + \Theta(e_2) + \Theta(e_3) < \pi$.

(∠5) Suppose that $G = (V, E)$ has an essentially unique mostly-embedding $\psi : G \to \mathbb{S}^2$. Let $e_1 = \langle u_1, v_1 \rangle, e_2 = \langle u_2, v_2 \rangle \in E$ so that $\psi(e_1)$ and $\psi(e_2)$ cross. Then u_1, u_2, v_1, v_2 form a complete subgraph of G. Let $\Theta : E \to [0, \pi)$. Then we require that $\Theta(e_1) = \Theta(e_2) = 0$, and that $\sum_{e \neq e_1, e_2} \Theta(e) = 2\pi$, with the sum taken over all other edges $e \neq e_1, e_2$ in the subgraph formed by u_1, u_2, v_1, v_2.

Conditions (∠4) and (∠5) can be thought of as the converses of one another. Condition (∠4) is necessary because if a 4-cycle in G formed by edges e_1, e_2, e_3, e_4 satisfies $\sum_{i=1}^{4} \Theta(e_i) = 2\pi$, then by a picture similar to that given in Figure 2.5 there is no "hole" between the disks corresponding to the endpoints of the e_i, and furthermore by thinness the disks must in fact meet as in Figure 2.4. Similarly, condition (∠5) is necessary for existence of a thin disk configuration by Observation 2.30.

Unfortunately, these conditions on (G, Θ) are not sufficient for existence of a thin disk configuration realizing (G, Θ), as we shall see in Section 2.9.6. However, we will prove the following:

Theorem 2.44 (Sufficient conditions for existence of a thin disk configuration)**.** *Let $G = (V, E)$ be a mostly-triangulation of the 2-sphere $\mathbb{S}^2$, and let $\Theta : E \to [0, \pi)$ so that conditions (∠1), (∠2), (∠3), (∠4), (∠5) all hold. Suppose also that the following condition holds:*

- *For every pair of successive edges $e_1 = \langle v_1, v_2 \rangle, e_2 = \langle v_2, v_3 \rangle$ in E, with $v_1 \neq v_3$, we have that $\Theta(e_1) + \Theta(e_2) \leq \pi$.*

Then there exists a disk configuration $\mathcal{C}$ on the Riemann sphere $\hat{\mathbb{C}}$ realizing (G, Θ), and it is thin. Furthermore $\mathcal{C}$ is uniquely determined by (G, Θ), up to action by Möbius and anti-Möbius transformations.

We prove this theorem in Section 2.9.5. The example in Section 2.9.6 will show that the extra hypothesis on (G, Θ) given in Theorem 2.44 is not necessary for the existence of a thin disk configuration realizing (G, Θ).

2.9.3 Hyperbolic polyhedra and Rivin's Theorem 2.45

We digress briefly to recall some ideas from hyperbolic geometry. A theorem in this area will be used to give us an existence statement in our setting. We will describe the connection between this setting and ours in Section 2.9.4.

We work in hyperbolic 3-space $\mathbb{H}^3$. A *hyperbolic half-space* of $\mathbb{H}^3$ is a subset of $\mathbb{H}^3$ the boundary of which is a hyperbolic plane. Then a *hyperbolic polyhedron* is the intersection of finitely many closed hyperbolic half-spaces, so that the resulting object has finite hyperbolic volume and non-empty interior. Such an object may alternatively be described as the closure of $\mathbb{H}^3 \setminus \cup_{i=1}^n H_i$, where each H_i is a hyperbolic half-space, again insisting that what remains has finite hyperbolic volume and non-empty interior.

Identify $\hat{\mathbb{C}}$ with the unit sphere in $\mathbb{R}^3$, and identify hyperbolic 3-space $\mathbb{H}^3$ with the open unit ball in $\mathbb{R}^3$ via the Poincaré ball model. In this model a hyperbolic plane is precisely the intersection of $\mathbb{H}^3$ with a Euclidean sphere or plane in $\mathbb{R}^3$ that meets $\hat{\mathbb{C}} = \partial\mathbb{H}^3$ orthogonally. Thus for example cutting out a hyperbolic half-space H_f from $\mathbb{H}^3$ may be visualized as taking an "ice cream scoop" out of $\mathbb{H}^3$.

Recall also that a *combinatorial polyhedron* X is a polytopal decomposition of the 2-sphere $\mathbb{S}^2$. The boundary of a hyperbolic polyhedron P decomposes naturally into vertices, edges, and faces, giving us a combinatorial polyhedron X which we say is *realized* by P. We adopt the convention that we will use the same symbol to refer to a vertex, edge, or face of P and the respective one of X, and trust that context will always make our meaning clear.

Recall that under the Poincaré ball model, the Riemann sphere $\hat{\mathbb{C}}$ can be thought of as $\partial\mathbb{H}^3$, or the "infinity" of $\mathbb{H}^3$. Then isometries of $\mathbb{H}^3$ are identified in the natural way with Möbius and anti-Möbius transformations of $\hat{\mathbb{C}}$.

An *ideal hyperbolic polyhedron* is a hyperbolic polyhedron all of whose vertices lie on $\hat{\mathbb{C}}$. The following theorem is a complete characterization of ideal hyperbolic polyhedra:

Rivin's Theorem 2.45. *Let $X = (V, E, F)$ be a combinatorial polyhedron, and let $\Theta : E \to (0, \pi)$ be a function. Then there exists an ideal hyperbolic polyhedron P in $\mathbb{H}^3$ combinatorially equivalent to X, so that the exterior dihedral angle at every edge e is given by $\Theta(e)$, if and only if the following two conditions hold:*

- *If v is a vertex of X then we have $\sum_{e \ni v} \Theta(e) = 2\pi$.*
- *Let $E_0 \subset E$ be a non-trivial cut-set of edges of X, that is, a set of edges of X so that $\cup_{e \in E_0} \text{interior}(e)$ separates the graph (V, E), but so that there is no vertex v belonging to every edge of E_0. Then $\sum_{e \in E_0} \Theta(e) > 2\pi$.*

Furthermore P, if it exists, is uniquely determined by X and Θ up to action by hyperbolic isometries.

We define the *exterior dihedral angle* between two faces of P meeting along an edge to be π minus the angle made between the faces inside of P. This theorem is due to Rivin. He obtains existence in [Riv96, Theorem 0.1], uniqueness in [Riv94, Theorem 14.1] and unifies and generalizes these results in [Riv03].

2.9.4 The relationship between hyperbolic polyhedra and disk configurations

Disk configurations on $\hat{\mathbb{C}}$ are closely related to polyhedra in hyperbolic 3-space $\mathbb{H}^3$. This relationship comes from the observation that a hyperbolic polyhedron having faces indexed by the set F may be described as the closure of $\mathbb{H}^3 \setminus \cup_{f \in F} H_f$, where every H_f is a hyperbolic half-space. Then under our identification of $\mathbb{H}^3$ with the unit ball in $\mathbb{R}^3$ via the Poincaré ball model, the boundaries of the H_f at infinity $\hat{\mathbb{C}} = \partial\mathbb{H}^3$ are metric closed disks in $\hat{\mathbb{C}}$. Denote by D_f the closed disk corresponding to H_f in this way. Then:

Observation 2.46. *If two faces of P, which meet along an edge, lie in ∂H_{f_1} and ∂H_{f_2} respectively, then their exterior dihedral angle is equal to $\pi - \measuredangle(D_{f_1}, D_{f_2})$.*

Let X be a combinatorial polyhedron. Then the *Poincaré dual*, or just *dual*, of X, which we denote X^*, is obtained by placing a dual vertex inside of each face of X, and connecting dual vertices by an edge in X^* if and only if the corresponding faces meet along an edge in X.

Suppose P is a hyperbolic polyhedron, having the combinatorics of the combinatorial polyhedron X. Then there is a disk configuration $\mathcal{C}$ naturally associated to $\mathcal{P}$ as described two paragraphs ago. Let G denote the contact graph of $\mathcal{C}$. We have the

following bijective associations:

$$\begin{array}{ccccc} \text{faces of } X & \longleftrightarrow & \text{faces of } P & \longleftrightarrow & \text{disks of } \mathcal{C} \\ \updownarrow & & & & \updownarrow \\ \text{vertices of } X^* & \dashleftarrow & \cdots\cdots & \dashrightarrow & \text{vertices of } G \end{array}$$

Furthermore, two vertices in X^* share an edge if and only if the associated faces of X share an edge, in which case the associated disks in $\mathcal{C}$ overlap, and so then there is also an edge between the associated vertices in G. Thus the 1-skeleton of X^* is a sub-graph of G:

$$\text{edges of the 1-skeleton of } X^* \hookrightarrow \text{edges of } G$$

Unfortunately, it does not turn out to be true that G and the 1-skeleton of X^* are equal: there may be extra contacts in $\mathcal{C}$ which are not reflected in the combinatorial polyhedra X and X^*. We shall see an example of this in Section 2.9.6. Nevertheless, theorems on hyperbolic polyhedra can be used to recover theorems on disk configurations. In particular, in Section 2.9.5, we will apply Rivin's Theorem 2.45 to give a proof of Theorem 2.44, which gives a partial answer to Question 2.43.

2.9.5 Proof of Theorem 2.44

The following may be proved from Rivin's Theorem 2.45:

Proposition 2.47. *Let $G = (V, E)$ be a mostly-triangulation of the 2-sphere $\mathbb{S}^2$, and let $\Theta : E \to [0, \pi)$ so that conditions ($\measuredangle 1$), ($\measuredangle 2$), ($\measuredangle 3$), ($\measuredangle 4$), ($\measuredangle 5$) all hold. Then there exists a disk configuration $\mathcal{C}$ on the Riemann sphere $\hat{\mathbb{C}}$ so that the following hold:*

- *There is a disk in $\mathcal{C}$ for every vertex of G.*
- *The contact graph of $\mathcal{C}$ may not equal G, but has G as a sub-graph.*
- *If $e = \langle u, v \rangle$ is an edge of G then $\measuredangle(D_u, D_v) = \Theta(e)$.*

Furthermore $\mathcal{C}$ is uniquely determined by (G, Θ), up to action by Möbius and anti-Möbius transformations. Also:

- *Every extraneous contact between two disks in $\mathcal{C}$ occurs inside of a third. That is, if $D_u, D_v \in \mathcal{C}$ meet, but $\langle u, v \rangle$ is not an edge of G, then $D_u \cap D_v$ is contained in some other disk of $\mathcal{C}$.*

- *The configuration $\mathcal{C}$ may not be thin, but the only places where the thinness condition fails are at these extraneous contacts if such exist.*

We sketch the construction required, as it is not completely trivial. An example is worked out in more detail in Section 2.9.6.

Proof sketch. Let $\psi : G \to \mathbb{S}^2$ be a mostly-embedding of G. We construct a combinatorial polyhedron from our mostly-embedding of G in the following way:

- Let $e_1 = \langle v_2, v_3 \rangle, e_2 = \langle v_1, v_3 \rangle, e_3 = \langle v_1, v_2 \rangle$ be three edges of G, so that their images $\psi(e_i)$ bound a triangle, call it f, in $\mathbb{S}^2$, so that the interior of f does not contain any of the embedded vertices of G. If $\sum_{i=1}^{3} \Theta(e_i) = \pi$, then do nothing. However, if $\sum_{i=1}^{3} \Theta(e_i) < \pi$, then add a vertex $\hat{v}$ inside of f, and embed new edges $\hat{e}_i = \langle \hat{v}, v_i \rangle$, $i = 1, 2, 3$ inside of f. This procedure subdivides f into three smaller triangles. Denote by $\hat{f}_1$ the sub-triangle of f bounded by $e_1, \hat{e}_2, \hat{e}_3$, similarly $\hat{f}_2, \hat{f}_3$. Add the edges $\hat{e}_i$ to our embedded drawing of G. It is not hard to check that there is a unique way to extend Θ to be defined on these new edges $\hat{e}_i$ so that $\sum_{e \subset \hat{f}_i} \Theta(e) = \pi$ for $i = 1, 2, 3$. This procedure may be thought of as adding a vertex which will correspond to a "dual disk."

- For any edge $e \in E$ so that $\Theta(e) = 0$, delete e from our embedded drawing. One may check, and we note, that no edge $\hat{e}$ added in the first step will have $\Theta(\hat{e}) = 0$. Note also that if $e_1 = \langle u_1, v_1 \rangle, e_2 = \langle u_2, v_2 \rangle \in E$ are edges of G having images $\psi(e_1)$ and $\psi(e_2)$ which cross under the mostly-embedding ψ, then necessarily $\Theta(e_1) = \Theta(e_2) = 0$, so we delete these e_1, e_2 from our drawing.

What remains after these two steps is an embedding in $\mathbb{S}^2$ of some planar graph, call it G_{X^*}, which divides $\mathbb{S}^2$ into triangles and quadrilaterals. Let X^* be the resulting combinatorial polyhedron. Then in particular G_{X^*} is the 1-skeleton of X^*. Let X be the combinatorial polyhedron obtained as the Poincaré dual of X^*.

Let $\Theta_{G_{X^*}}$ be defined on the edges of G_{X^*} either by restriction of Θ or as described in the construction of X, as necessary. Let Θ_X be defined on the edges of X in the following way: using the same symbol e to refer to an edge of X and its associated edge in G_{X^*}, set $\Theta_X(e) = \pi - \Theta_{G_{X^*}}(e)$. We do this because, as per Observation 2.46, the exterior dihedral angle between two faces of a hyperbolic polyhedron is π minus the exterior intersection angle of the associated closed disks in $\hat{\mathbb{C}}$.

It is straightforward to check that the pair (X, Θ_X) satisfies the hypotheses of Rivin's Theorem 2.45. Thus, let P be the ideal hyperbolic polyhedron for (X, Θ_X) obtained from Rivin's Theorem 2.45. Let $\mathcal{C}_0$ be the associated disk configuration.

Finally, obtain $\mathcal{C}$ from $\mathcal{C}_0$ by deleting those disks corresponding to the vertices which were added to X during its construction.

Using this construction, it is mostly routine to check the statement of Proposition 2.47. The following lemma helps:

Lemma 2.48. *Let P be an ideal hyperbolic polyhedron combinatorially equivalent to the combinatorial polyhedron $X = (V, E, F)$. Write P as the closure of $\mathbb{H}^3 \setminus \cup_{f \in F} H_f$, where every H_f is a hyperbolic half-space, so that H_f contains in its boundary hyperplane the face corresponding to f in P. For $f \in F$ write $D_f \subset \hat{\mathbb{C}} = \partial\mathbb{H}^3$ to denote the boundary disk at infinity of H_f. Suppose that $g, h \in F$ are faces of X so that the associated disks D_g and D_h meet, but so that g and h do not share an edge in X.*

Then there is a finite sequence of faces $g = f_1, f_2, \ldots, f_n = h$ so that for all $1 \leq i < n$,

- *the faces f_i and f_{i+1} share an edge in X, and*
- *the intersection $D_g \cap D_h$ is contained in $D_{f_i} \cap D_{f_{i+1}}$.*

Furthermore, for $1 \leq i < n-1$, we have $\measuredangle(D_i, D_{i+1}) + \measuredangle(D_{i+1}, D_{i+2}) > \pi$.

We omit the proof of Lemma 2.48. Theorem 2.44 follows from Proposition 2.47 and Lemma 2.48.

We remark that the rigidity portion of Proposition 2.47 follows from the rigidity in Rivin's Theorem 2.45. Our own rigidity proofs do not apply, because as we say in the statement, the disk configurations resulting from Proposition 2.47 are not thin. Our Main Rigidity Theorem 1 does imply the rigidity portion of Theorem 2.44, but this is already guaranteed by the rigidity of Rivin's Theorem 2.45.

2.9.6 Examples of combinatorial mismatching between hyperbolic polyhedra and their associated disk configurations

In this section, we construct examples illustrating the following phenomena:

- There exist two ideal hyperbolic polyhedra P and $\tilde{P}$, which realize the same combinatorial polyhedron, but so that the following holds: if we denote by $\mathcal{C}$ and $\tilde{\mathcal{C}}$ the disk configurations corresponding to P and $\tilde{P}$ respectively, then the contact graphs of $\mathcal{C}$ and $\tilde{\mathcal{C}}$ do not agree. This is Example 2.49.

- There exists a graph $G = (V, E)$ and $\Theta : E \to [0, \pi)$, so that (G, Θ) satisfies conditions $(\measuredangle 1)$, $(\measuredangle 2)$, $(\measuredangle 3)$, $(\measuredangle 4)$, $(\measuredangle 5)$ non-vacuously, and so that the following holds: if $\mathcal{C}'$ is any disk configuration whose incidence data has (G, Θ) as sub-incidence-data, then there is at least one extra edge connecting two of the vertices of the canonically embedded image of G in the contact graph of $\mathcal{C}'$. This is Example 2.52. Finally:

- There are examples of incidence data in which "extraneous contacts" may be forced by "non-local" conditions in the data. In particular, let v_1, v_2 be two vertices having an "extra contact," in our previous example. It turns out that this extra contact, between disks D_1 and D_2, is contained inside of a third disk $D_3 \in \mathcal{C}$. Let v_3 be the vertex corresponding to this disk D_3. Then, in our example, we may modify the values taken by Θ at edges which are "away" from v_1, v_2, v_3, that is, edges none of which have any of v_1, v_2, v_3 as an endpoint, so that the following holds: denote by $(\tilde{G}, \tilde{\Theta})$ the incidence data we are left with after our modification. Then there is a disk configuration having precisely $(\tilde{G}, \tilde{\Theta})$ as its incidence data. This is Example 2.54.

Example 2.49. Combinatorially equivalent ideal hyperbolic polyhedra having combinatorially unequivalent associated disk configurations

We explicitly construct two ideal hyperbolic polyhedra P and $\tilde{P}$ realizing the same combinatorial polyhedron, in the following way:

- We begin with two explicitly chosen disk configurations $\mathcal{C}$ and $\tilde{\mathcal{C}}$, having unequal contact graphs. These are pictured in Figure 2.6.

- We add "dual disks" to the configurations to obtain configurations, which we denote $\mathcal{C}_P$ and $\tilde{\mathcal{C}}_{\tilde{P}}$, each of which completely covers $\hat{\mathbb{C}}$. These are the configurations whose disks will cut out our ideal hyperbolic polyhedra.

- We define P as the closure of $\mathbb{H}^3 \setminus \cup_{D \in \mathcal{C}_P} H_D$, where H_D is the hyperbolic half-space corresponding to the closed disk $D \subset \hat{\mathbb{C}} = \partial \mathbb{H}^3$, defining $\tilde{P}$ similarly.

We then argue that the ideal hyperbolic polyhedra P and $\tilde{P}$ constructed in this way from our explicitly chosen $\mathcal{C}$ and $\tilde{\mathcal{C}}$ realize the same combinatorial polyhedron X. As the punchline of this example, we will observe that the configurations $\mathcal{C}_P$ and $\tilde{\mathcal{C}}_P$ have unequal contact graphs, as do $\mathcal{C}$ and $\tilde{\mathcal{C}}$. Our starting configurations $\mathcal{C}$ and $\tilde{\mathcal{C}}$ are shown in Figure 2.6.

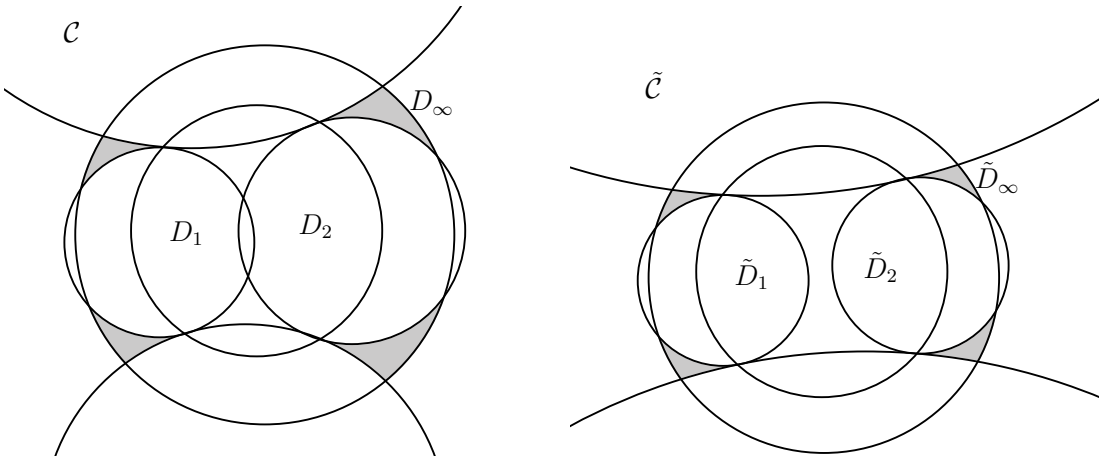

Figure 2.6: **Two disk configurations having unequal contact graphs.** The contact graphs are unequal because of the "extra" contact between D_1 and D_2 in $\mathcal{C}$. Here the disks labeled D_∞ and $\tilde{D}_\infty$ are centered at $\infty \in \hat{\mathbb{C}}$, and all other disks are contained in $\hat{\mathbb{C}} \setminus \{\infty\}$. The shaded regions are $\hat{\mathbb{C}} \setminus \cup_{D \in \mathcal{C}} D$ and $\hat{\mathbb{C}} \setminus \cup_{\tilde{D} \in \tilde{\mathcal{C}}} \tilde{D}$.

We now add our "dual disks." First, note that $\hat{\mathbb{C}} \setminus \cup_{D \in \mathcal{C}} D$ consists of four curvilinear triangles. Let $\mathcal{C}_P$ be obtained from $\mathcal{C}$ in the following way: for each of these curvilinear triangle, add a closed disk covering this triangle, passing through its three corners. Construct $\tilde{\mathcal{C}}_{\tilde{P}}$ in the analogous way. The resulting configurations are shown in Figure 2.7. Note that there is one and only one way to carry out this portion of the construction from our starting configurations.

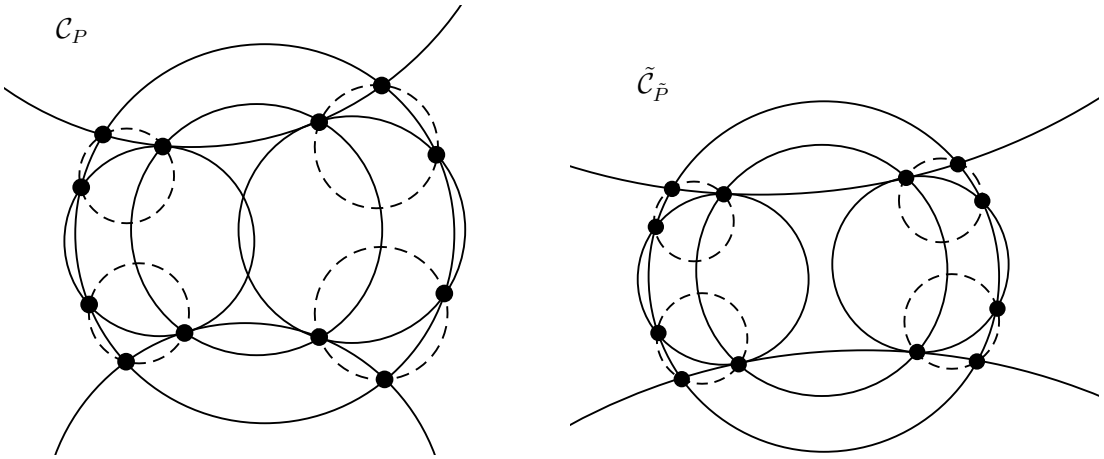

Figure 2.7: **The disk configurations of Figure 2.6, after adding "dual disks."** We have drawn the dual disks dashed. The marked vertices will turn out to be the vertices of the ideal hyperbolic polyhedra we eventually construct.

We define P to be the closure of $\mathbb{H}^3 \setminus \cup_{D \in \mathcal{C}_P} H_D$, where H_D is the hyperbolic half-space whose limit set at infinity $\hat{\mathbb{C}} = \partial \mathbb{H}^3$ is the closed disk D. Intuitively, we

"scoop out" hyperbolic half-spaces for all of the disks of $\mathcal{C}_P$. The disk D_∞ should be considered as a special case: it is centered at $\infty \in \hat{\mathbb{C}}$, so "scooping out" around D_∞ really looks like keeping only the "scoop" corresponding to the closed disk in $\mathbb{C}$ bounded by ∂D_∞. We define $\tilde{P}$ analogously, from $\tilde{\mathcal{C}}_{\tilde{P}}$. Our goal is to argue that P and $\tilde{P}$ are ideal hyperbolic polyhedra, and furthermore that they realize the same combinatorial polyhedron. In order to do so, we digress briefly to recall some facts from hyperbolic geometry.

For a moment, consider the upper half-space model of $\mathbb{H}^3$, identifying an isometrically embedded copy of the complex plane $\mathbb{C} \subset \mathbb{R}^3$ with $\partial\mathbb{H}^3$ minus a point. Then hyperbolic hyperplanes are obtained precisely by taking intersections of $\mathbb{H}^3$ with Euclidean spheres and planes meeting our embedded copy of $\mathbb{C} \subset \mathbb{R}^3$ orthogonally. Now, the intersection of two hyperplanes in $\mathbb{H}^3$ is a hyperbolic geodesic, and hyperbolic geodesics in this model are obtained precisely by taking intersections of $\mathbb{H}^3$ with Euclidean lines and circles which pass through $\mathbb{C} = \partial\mathbb{H}^3 \setminus \{\infty\}$ orthogonally. Thus the orthogonal projection of a hyperbolic geodesic to $\mathbb{C}$ in this model is a straight line segment or a single point. This discussion makes the following an easy observation:

Observation 2.50. *Suppose A, B, C are metric closed disks in $\mathbb{C}$, and let H_A, H_B, H_C respectively be the associated hyperbolic half-spaces. Suppose that $A \cap B \subset C$. Then $H_A \cap H_B \subset H_C$.*

Thus, under the hypotheses of Observation 2.50, if we "scoop out" H_A and H_B, we initially form an edge e between their boundary hyperplanes, but after "scooping out" H_C as well, this edge e will disappear, other than potentially one or both of its endpoints at infinity $\hat{\mathbb{C}} = \partial\mathbb{H}^3$.

We return to our setting, again considering $\mathcal{C}$ and $\tilde{\mathcal{C}}$. *A priori*, any two overlapping disks of $\mathcal{C}$ may contribute an edge to P, and any point which is an intersection point of several boundary circles of $\mathcal{C}$ may contribute a vertex to P. However, with Observation 2.50 in mind, one can show that the pairs $D_i, D_j \in \mathcal{C}$ of overlapping disks which contribute an edge to P are precisely those pairs whose intersection $D_i \cap D_j$ is not contained in any other disk of $\mathcal{C}$. Thus we may go through the drawings in Figure 2.7 and explicitly describe which overlaps will contribute edges to our polyhedra. Furthermore a vertex of an ideal hyperbolic polyhedron is exactly a common endpoint of several of its edges. Then it is easy to check that the vertices of P and $\tilde{P}$ are precisely the points labeled in Figure 2.7. Finally, it is routine to show that P and $\tilde{P}$ have finite hyperbolic volume via standard integral computations.

Then it is not hard to write down explicitly the combinatorial polyhedron realized by P, and that realized by $\tilde{P}$, and to see that they are the same. We show the combinatorial polyhedron X which P and $\tilde{P}$ both realize, as well as its dual X^*, in Figure 2.8.

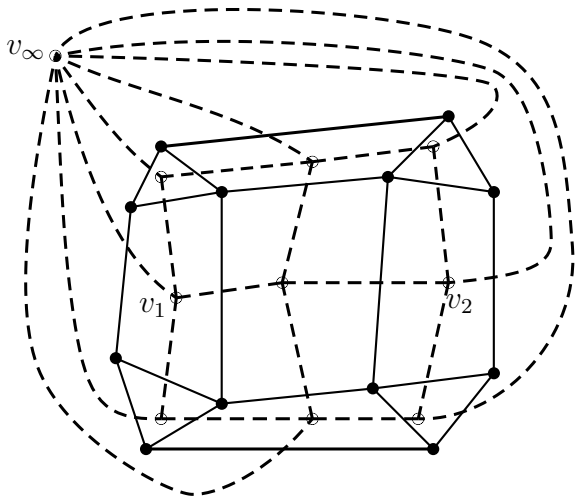

Figure 2.8: **The combinatorial polyhedron X realized by P, and its dual X^*.** We have drawn X with solid vertices and edges, and X^* with open vertices and dashed edges.

Let $X = (V, E, F)$ be the combinatorial polyhedron realized by P and $\tilde{P}$, and let $\Theta_P : E \to (0, \pi)$ and $\Theta_{\tilde{P}} : E \to (0, \pi)$ be obtained from the exterior dihedral angles of the respective edges of P and of $\tilde{P}$. Of course, in light of the uniqueness portion of Rivin's Theorem 2.45, we must have that Θ_P and $\Theta_{\tilde{P}}$ are not equal: if they were equal, then P and $\tilde{P}$ would differ by a hyperbolic isometry, which would give a Möbius or anti-Möbius transformation sending $\mathcal{C}$ to $\tilde{\mathcal{C}}$, but there can be no such transformation because such a transformation on disk configurations must be contact-graph-preserving. The following is a natural question:

Question 2.51. *Suppose that $X = (V, E, F)$ is a combinatorial polyhedron, and $\Theta : E \to (0, \pi)$, so that the pair (X, Θ) satisfy the hypotheses of Rivin's Theorem 2.45. Let P be an ideal hyperbolic polyhedron realizing (X, Θ). Write P as the closure of $\mathbb{H}^3 \setminus_{f \in F} H_f$, where each H_f is a hyperbolic half-space, whose boundary hyperplane contains the face f of P, and so that H_f limits to the disk D_f at infinity $\hat{\mathbb{C}} = \partial\mathbb{H}^3$. Suppose that f_1, f_2 are distinct faces of P sharing neither a vertex nor an edge. How can we read off, from the combinatorial information of (X, Θ), whether the associated disks D_{f_1}, D_{f_2} intersect?*

Certainly it is true that if f_1, f_2 are faces of X sharing a vertex but no edge, then the associated disks D_1, D_2 meet. A complete answer to Question 2.51 would give some partial answers to Question 2.41.

Example 2.52. Incidence data forcing additional contacts

We now describe incidence data (G, Θ) satisfying conditions $(\measuredangle 1)$, $(\measuredangle 2)$, $(\measuredangle 3)$, $(\measuredangle 4)$, $(\measuredangle 5)$ non-vacuously, so that the following holds: if $\mathcal{C}'$ is any disk configuration whose incidence data has (G, Θ) as sub-incidence-data, then there are extra edges among the vertices of the canonically embedded image of G in the contact graph of $\mathcal{C}'$.

Consider again the disk configuration $\mathcal{C}$ on $\hat{\mathbb{C}}$ shown in Figure 2.6. Let $(G_{\mathcal{C}} = (V_{\mathcal{C}}, E_{\mathcal{C}}), \Theta_{\mathcal{C}})$ be the incidence data which $\mathcal{C}$ realizes. Let $G = (V, E)$ be obtained from $G_{\mathcal{C}}$ by deleting the open edge $\langle v_1, v_2 \rangle$ representing the contact between D_1 and D_2. Let $\Theta : E \to [0, \pi)$ be obtained by restricting $\Theta_{\mathcal{C}}$ to E. For reference, we draw the graphs $G_{\mathcal{C}}$ and G, in Figure 2.9. Note that (G, Θ) satisfies each of our conditions $(\measuredangle 1)$, $(\measuredangle 2)$, $(\measuredangle 3)$, $(\measuredangle 4)$, $(\measuredangle 5)$ non-vacuously, in particular that G has a essentially unique embedding in $\mathbb{S}^2$. Note also that $\mathcal{C}$ is not thin.

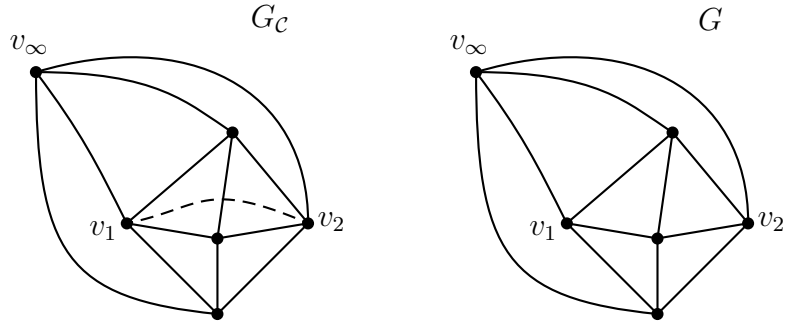

Figure 2.9: **The contact graph of $\mathcal{C}$, and the graph G obtained by removing an edge.** The vertices v_1, v_2, v_∞ correspond to the disks D_1, D_2, D_∞, respectively. The dashed edge in $G_{\mathcal{C}}$ represents the contact between D_1 and D_2.

We will see that if $\mathcal{C}'$ is a disk configuration realizing incidence data (G', Θ'), so that (G, Θ) is sub-incidence-data of (G', Θ'), then there is an "extraneous contact" between the disks corresponding to the vertices v_1 and v_2 in $\mathcal{C}'$. First, throw out any disks of $\mathcal{C}'$ which do not correspond to vertices of G' hit by the injection of the incidence data (G, Θ), and re-compute (G', Θ') to reflect this modification of $\mathcal{C}'$. Suppose for contradiction that $(G', \Theta') = (G, \Theta)$. If this were the case, then $\mathcal{C}'$ would be a disk configuration realizing (G, Θ) with no "extraneous contacts" among any of its disks, an event we hope to show is impossible.

Our strategy is to show that we may run the same construction on $\mathcal{C}'$ that we ran on $\mathcal{C}$, obtaining an ideal hyperbolic polyhedron P' combinatorially equivalent to P, and with the same exterior dihedral angle function. Then P and P' are hyperbolically isometric by Rivin's Theorem 2.45. The hyperbolic isometry sending P to P' is associated to a Möbius transformation that must send the overlapping disks $D_1, D_2 \in \mathcal{C}$ to the supposedly disjoint disks $D_1', D_2' \in \mathcal{C}'$, giving us our contradiction.

First, note that under our contradiction hypotheses we have that $\mathcal{C}'$ is thin. It follows without much trouble that if we draw every edge $e = \langle u, v \rangle$ of the contact graph G' of $\mathcal{C}'$ by connecting the spherical centers of the disks D_u, D_v corresponding to the endpoints of e with the spherical geodesic arc passing through $D_u \cap D_v$, we get an embedding ψ' of G' in $\mathbb{S}^2$. Note that because G' is the 1-skeleton of a triangulation of $\mathbb{S}^2$, it has an essentially unique geodesic embedding.

We next wish to establish that $\hat{\mathbb{C}} \setminus \cup_{D' \in \mathcal{C}'} D'$ consists of four curvilinear triangles, which at their corners have the same angles as the associated curvilinear triangles in $\hat{\mathbb{C}} \setminus \cup_{D \in \mathcal{C}} D$. Let $T' = (V', E', F')$ be the triangulation of $\mathbb{S}^2$ induced by the embedding ψ' of $G' = (V', E')$ described in the previous paragraph. This triangulation consists of 8 triangular faces, each of which may *a priori* contain at most one component of $\hat{\mathbb{C}} \setminus \cup_{D' \in \mathcal{C}'} D'$. However, observe the following:

Observation 2.53. *Suppose that $f' \in F'$ is a face of T', bounded by the edges e_1', e_2', e_2', so that $\sum_{i=1}^{3} \Theta'(e_i') = \pi$. Then the embedded image of f' under ψ' is completely covered by $D_1' \cup D_2' \cup D_3'$. Conversely, if $\sum_{i=1}^{3} \Theta'(e_i') < \pi$, then the embedded image of f' under ψ' contains a curvilinear triangle $U' \subset \hat{\mathbb{C}}$, bounded by circular arcs, so that the interior angles at the corners of U' are $\Theta(e_i')$.*

This observation gives us exactly what we want: keeping in mind that the incidence data of $\mathcal{C}$ restricts to that of $\mathcal{C}'$, inspection of $\mathcal{C}$ yields that there are precisely four faces f' of T' which satisfy $\sum_{e_i' \subset f'} \Theta'(e_i') = \pi$, and four which satisfy $\sum_{e_i' \subset f'} \Theta'(e_i') < \pi$.

At this point it is clear how to proceed: every component U' of $\hat{\mathbb{C}} \setminus \cup_{D' \in \mathcal{C}'} D'$ corresponds to a component U of $\hat{\mathbb{C}} \setminus \cup_{D \in \mathcal{C}} D$, and furthermore it is not hard to show that U and U' are Möbius equivalent. Thus we may add "dual disks" to $\mathcal{C}'$ as we did to $\mathcal{C}$, and obtain an ideal hyperbolic polyhedron P' via the construction analogous to our construction of P. It is not hard to verify that P and P' are combinatorially equivalent and that their exterior dihedral angle functions agree using our observations up to this point. We leave these last details to the reader.

Example 2.54. Eliminating extraneous contacts "non-locally"

In this last of our series of examples exploring the existence question for disk configurations, we see that we may

- take incidence data (G, Θ), the realization of which forces "extra contacts," say between the disks corresponding to vertices v_1, v_2 of G,
- modify Θ "far away" from v_1 and v_2, that is, modify the values taken by Θ only at edges which are not incident to one of v_1 and v_2, and
- thereby obtain incidence data that is realizable without "extra contacts."

This shows that a complete answer to Question 2.41 may be quite difficult, as extraneous contacts may be forced by incidence data which is in some sense localized "far away" from where these contacts happen.

We again consider the configurations $\mathcal{C}$ and $\tilde{\mathcal{C}}$ from Figure 2.6. Let $(G_{\mathcal{C}}, \Theta_{\mathcal{C}})$ and $(G_{\tilde{\mathcal{C}}}, \Theta_{\tilde{\mathcal{C}}})$ be the incidence data realized by $\mathcal{C}$ and $\tilde{\mathcal{C}}$ respectively. Let (G, Θ) be obtained from the incidence data $(G_{\mathcal{C}}, \Theta_{\mathcal{C}})$ by deleting the edge $\langle v_1, v_2 \rangle$ as before.

Note that then G and $G_{\tilde{\mathcal{C}}}$ are equal. Identify their edges in the natural way. It is not hard to see that it is possible to construct $\mathcal{C}$ and $\tilde{\mathcal{C}}$ explicitly, keeping our discussion and figures up to this point qualitatively the same, so that Θ and $\Theta_{\tilde{\mathcal{C}}}$ agree on every edge having v_1 or v_2 as an endpoint. However, we have already shown that there can be no realization of (G, Θ), and $\tilde{\mathcal{C}}$ is an explicit realization of $(G_{\tilde{\mathcal{C}}}, \Theta_{\tilde{\mathcal{C}}})$, completing our example.

2.9.7 Closing remarks on the existence question

Absent the thinness condition, if we hoped to get existence statements in $\hat{\mathbb{C}}$ similar to Theorem 2.44 from the setting of hyperbolic polyhedra, we would require existence statements on non-ideal polyhedra in $\mathbb{H}^3$, that is, hyperbolic polyhedra some of whose vertices may lie in the interior of $\mathbb{H}^3$. Results along these lines have been obtained by Rivin and others, see for example [HR93, Theorem 1.1]. The idea to study circle packings and disk configurations via hyperbolic polyhedra dates back to Thurston. We discuss this further in Section 2.10.

Given an existence theorem for thin disk configurations on $\hat{\mathbb{C}}$, it seems likely that existence theorems for configurations in $\mathbb{C}$ and $\mathbb{D} \cong \mathbb{H}^2$ would follow, via a construction analogous to the one used to prove the existence portion of the Discrete Uniformization Theorem 2.15, using the Circle Packing Theorem 2.2 and the Koebe–Andreev–Thurston Theorem 2.8, as outlined in Section 2.4. We would require gen-

eralizations of the Discrete Ring Lemma 2.21 and the Discrete Schwarz–Pick Lemma 2.19 for the construction portion of the proof. Fortunately, statements along those lines are proved in [He99, Lemmas 2.2, 7.1]. Next, we would need to generalize the proof of local finiteness of the limiting object in our construction. The steps of the topological argument given by Stephenson in [Ste05, Section 8.2] seem to generalize cleanly, except that at one point he appeals to a fixed-point index argument to prove a theorem on "discrete harmonic measure." Our main technical result, the Index Theorem 4, would exactly give the required statement on fixed-point index to generalize this last remaining portion of the argument. Approaches to the existence question via vertex extremal length arguments may be possible as well, and have been carried out in similar settings by Z.-X. He. We will discuss this further in Section 2.10.

Thanks to Lars Louder for fruitful discussions related to the issues arising Figure 2.6. Thanks to my advisor Jeff Lagarias for fruitful discussions on the existence question for disk configurations.

2.10 Related results

We are not the first to consider configurations of overlapping disks. In fact, consideration of such configurations dates back to Thurston. We state the result most similar to ours currently present in the literature, and then discuss the timeline of the proofs of its various parts and special cases.

Theorem 2.55. *Let $G = (V, E)$ be a graph which mostly-triangulates an oriented topological surface S, with mostly-embedding $\psi : G \to S$. Let $\Theta : E \to [0, \pi/2]$ be so that the following conditions hold:*

(2.55a) Let $e_1 = \langle u_1, v_1 \rangle, e_2 = \langle u_2, v_2 \rangle \in E$ be edges of G so that $\psi(e_1)$ and $\psi(e_2)$ cross. Then $\Theta(e_1) = \Theta(e_2) = 0$, and for any other edge e whose endpoints are two of u_1, u_2, v_1, v_2, we have $\Theta(e) = \pi/2$.

(2.55b) If the edges e_1, e_2, e_3, e_4 form a 4-cycle in (V, E) with $\Theta(e_i) = \pi/2$ for all i, then there is at least one more edge connecting the vertices of the e_i.

(2.55c) If the edges e_1, e_2, e_3 form a 3-cycle in (V, E) that separates the graph (V, E), then $\Theta(e_1) + \Theta(e_2) + \Theta(e_3) < \pi$.

Then there is a complete constant curvature Riemannian metric d on S and a disk configuration $\mathcal{C}$ locally finite in (S, d) realizing (G, Θ). Furthermore d is uniquely determined by (G, Θ) up to conformal isomorphisms, and $\mathcal{C}$ is then uniquely determined by (G, Θ) and d up to conformal and anti-conformal automorphisms of (S, d).

We first discuss the relationship between the disk configurations considered in Theorem 2.55 and our thin disk configurations. The key point is that under the constraint $\Theta \leq \pi/2$, we again get that the only way that four disks may meet at any point is the way shown in Figure 2.4 on p. 31. Furthermore in this case, if four disks meet at a point in that way, then the non-zero overlap angles among the four disks must be exactly $\pi/2$, hence Condition (2.55a). Thus apparently proofs in the setting of overlapping disks become tractable once we have control over the way that four or more disks of a configuration may intersect. In this thesis, we control these intersections via the thinness condition, while it seems that the historical approach has been to control these intersections via the condition that $\Theta \leq \pi/2$.

Next, we compare the conditions on (G, Θ) given in the statement of Theorem 2.55 to our conditions ($\measuredangle 1$), ($\measuredangle 2$), ($\measuredangle 3$), ($\measuredangle 4$), ($\measuredangle 5$), found on p. 36. First note that (2.55a) is equivalent to ($\measuredangle 5$) in the case when $\Theta \leq \pi/2$. Similarly, in the case where $\Theta \leq \pi/2$, condition (2.55b) follows from condition ($\measuredangle 4$). Next, we have that (2.55c) is exactly ($\measuredangle 3$). Finally, observe that ($\measuredangle 1$) is satisfied automatically in the case $n > 4$ if $\Theta \leq \pi/2$, and follows vacuously from (2.55b) for $n = 4$, again supposing $\Theta \leq \pi/2$. Thus Theorem 2.55 as a special case establishes existence and uniqueness of thin disk configurations whose incidence information (G, Θ) satisfies our conditions ($\measuredangle 1$), ($\measuredangle 2$), ($\measuredangle 3$), ($\measuredangle 4$), ($\measuredangle 5$), provided $\Theta \leq \pi/2$. It also establishes Conjecture 2.37 in this special case. We remark also that as a corollary of this discussion, we see that conditions (2.55a), (2.55b), (2.55c) are not only sufficient for existence in the statement of Theorem 2.55, but necessary.

We now move on to our discussion of the timeline of the development of Theorem 2.55. Thurston was aware of both the existence and the uniqueness parts in the case of the 2-sphere $S = \mathbb{S}^2$ in the 1980s, see comments in [Thu80, Section 13.6, p. 333]. His approach was via his interpretation of Andreev's Theorem, see [And70], in the style of our proof of Proposition 2.47. Andreev's Theorem gives a complete characterization of ideal hyperbolic polyhedra analogous to that given by Rivin's Theorem 2.45 in the special case when $\Theta \leq \pi/2$. Rivin remarks that in the setting of polyhedra, the restriction $\Theta \leq \pi/2$ is a very strong one, see comments in [Riv96, p. 52]. In particular, in the setting of his and Andreev's theorems, it implies that the valence of every vertex of the resulting ideal polyhedron is either 3 or 4.

Thurston also proved Theorem 2.55 directly for closed surfaces of positive genus, see [Thu80, Theorem 13.7.1]. The idea of the proof is to run an infinite constructive iterative algorithm, adjusting the radii of the disks one at a time, obtaining the desired disk configuration in the limit.

Next, Z.-X. He proved Theorem 2.55 in the case where S is a topological open disk in [He99]. His method of proof is analytical, employing vertex extremal length and the theory of electrical networks on graphs. In the same paper, he wrote that he was preparing a paper which would handle the case where Θ may exceed $\pi/2$, but he never published such a paper, and has since retired from academic mathematics. As usual, Theorem 2.55 in the case where S is simply connected implies the full statement via our argument using covering space theory.

Bobenko and Springborn have proved a theorem similar to Theorem 2.55 where they allow Θ to exceed $\pi/2$, considering only compact surfaces, in [BS04, Theorem 4]. In particular, they do not handle disk configurations which live naturally in $\mathbb{C}$ or $\mathbb{H}^2$. The style of their theorem most closely resembles Rivin's Theorem 2.45, and they call it the higher-genus analog of Rivin's Theorem 2.45. Their proof uses variational principles.

2.11 Dualizable packings

The work that eventually led to the results of this thesis began with the following question:

Question 2.56. *Suppose that we start with a polytopal decomposition of the sphere, instead of a triangulation. Can we recover a uniqueness statement in the style of the Koebe–Andreev–Thurston Theorem 2.8?*

Of course, the CPT 2.2 still applies, so there exist circle packings realizing the contact graph of any polytopal decomposition, but it is easy to construct packings realizing the same polytopal decomposition which are not Möbius equivalent.

The "right" way to answer Question 2.56 is to consider so-called dualizable circle packings. An *interstice* of a circle packing is a topological open disk U formed by the boundaries of some disks of the packing, which is minimal in the sense that no other disks of the packing meet U. An interstice has as its "corners" the points of tangency of the disks along its boundary. Then the circle packing $\mathcal{P} = \{D_v\}_{v \in V}$ is called *dualizable* if

- for every interstice U of $\mathcal{P}$ there is a disk D_U^* whose boundary circle passes through all of the "corners" of U, so that in addition,
- for every disk D_v along the boundary of the interstice U, the disks D_U^* and D_v meet orthogonally, meaning that their boundary circles have orthogonal tangent vectors at the two points where these circles meet.

For example, every circle packing realizing a triangulation is automatically dualizable. We will call D_U^* the *interstitial disk* for U. Then we get the following:

Dualizable Circle Packing Theorem 2.57. *Let X be a polytopal decomposition of the sphere. Then there exists a dualizable circle packing realizing X, and it is uniquely determined by X up to Möbius and anti-Möbius transformations.*

Note that if $\mathcal{P}$ is any circle packing realizing the polytopal decomposition $X = (V, E, F)$, then the interstices of $\mathcal{P}$ are in bijection with the faces F of the decomposition. Then suppose that $\mathcal{P}$ is dualizable, and that f is the face corresponding to the interstice U of $\mathcal{P}$. Denote $D_f^* = D_U$. Then $\mathcal{P}^* = \{D_f^*\}_{f \in F}$ is a dualizable circle packing realizing X^* the dual polytopal decomposition to X. The packing $\mathcal{P}^*$ is called the *dual packing* to $\mathcal{P}$.

The DCPT 2.57 has an unclear history. Vague comments in [Thu80] suggest that Thurston may have been aware of it as early as 1980, although it cannot be attributed to him, as he never gave the statement in print. Certainly Thurston's proof of the KATT 2.8 via the Mostow–Prasad Rigidity Theorem 2.58 indicates that he was at least aware of the notion of dual circle packings. We sketch this proof now, in part because it is a nice example of a deep connection between circle packing and three-dimensional geometry and topology. First, recall:

Mostow–Prasad Rigidity Theorem 2.58. *Let Γ and $\tilde{\Gamma}$ be discrete subgroups of the isometry group of hyperbolic 3-space $\mathbb{H}^3$. Suppose that both quotients $\mathbb{H}^3/\Gamma$ and $\mathbb{H}^3/\tilde{\Gamma}$ have finite volume, and suppose that Γ and $\tilde{\Gamma}$ are isomorphic as discrete groups. Then Γ and $\tilde{\Gamma}$ are conjugate.*

For a reference see [Thu80, Theorem 5.7.1]. Actually the MPRT 2.58 holds for dimension $n \geq 3$, but we need only the dimension 3 version. Then Thurston's proof proceeds as follows. Identify $\mathbb{H}^3$ with the open unit ball in $\mathbb{R}^3$ via the Poincaré ball model, so that $\partial\mathbb{H}^3$ is identified with the Riemann sphere $\hat{\mathbb{C}}$. Then isometries of $\mathbb{H}^3$ are identified with Möbius and anti-Möbius transformations of $\hat{\mathbb{C}}$. Suppose that two circle packings $\mathcal{P}$ and $\tilde{\mathcal{P}}$ realize the same triangulation of the sphere. We wish to show that the two are equivalent via Möbius or anti-Möbius transformations. Let Γ be the group generated by reflections in the boundary circles of the disks of $\mathcal{P}$ and those of the dual packing $\mathcal{P}^*$, similarly $\tilde{\Gamma}$ for $\tilde{\mathcal{P}}$ and $\tilde{\mathcal{P}}^*$. Because $\mathcal{P}$ and $\tilde{\mathcal{P}}$ have the same combinatorics, the discrete groups Γ and $\tilde{\Gamma}$ are isomorphic. Also, it is not hard to check that $\mathbb{H}^3/\Gamma$ and $\mathbb{H}^3/\tilde{\Gamma}$ have finite volume. Then $\mathcal{P}$ and $\tilde{\mathcal{P}}$ are equivalent under the Möbius transformation which conjugates Γ to $\tilde{\Gamma}$. Of course, this same proof establishes the rigidity portion of the DCPT 2.57.

The first published proof of the DCPT 2.57 is generally attributed to Brightwell and Scheinerman. They give a proof in [BS93]. Although other papers containing the statement or generalizations thereof predate theirs, these other papers directly reference their preprint. There, the proof is via an interpretation of Thurston's methods. The interpretation is attributed by the authors to Lovász. As an application, they answer the following question, posed in [Tut63, §13]:

Question 2.59 (Tutte)**.** *Given a finite planar graph $G = (V, E)$, is there a simultaneous straight-line embedding of it and its planar dual $G^* = (V^*, E^*)$, so that every edge $e \in E$ crosses its partner $e^* \in E^*$ at a right angle?*

The DCPT 2.57 gives the best such embedding possible. The DCPT 2.57 is also stated and discussed in [Sac94, Theorem 2], a survey article by Sachs. There he also gives further references to early papers addressing it.

The following can be proved a corollary of Theorem 2.55:

Dualizable Discrete Uniformization Theorem 2.60. *Let $X = (V, E, F)$ be a polytopal decomposition of an oriented topological surface S. Then there is a complete constant curvature Riemannian metric d on S and a locally finite dualizable circle packing $\mathcal{P}$ in (S, d) which realizes X. Furthermore d is uniquely determined by X up to conformal isomorphisms, and $\mathcal{P}$ is then uniquely determined by X and d up to conformal and anti-conformal automorphisms of (S, d).*

The idea of the proof of the DDUT 2.60 from Theorem 2.55 is as follows. Start with a polytopal decomposition X. Inside every face f, add a new vertex, and join it with an edge to every vertex on ∂f. Set Θ equal to $\pi/2$ for every new edge added this way, and set Θ equal to 0 for every original edge of X. Then apply Theorem 2.55. This is slightly technically imprecise, but it is clear how to proceed. Alternatively, our rigidity theorems accommodate this argument, so rigidity of dualizable circle packings is also a consequence of our main rigidity results.

Mohar proved Theorem 2.60 in the case where S is a closed surface in [Moh93], via an adaptation of Thurston's iterative algorithm for constructing circle packings. He includes a run-time analysis of his algorithm. Other than this, we are not aware of Theorem 2.60 or any special cases of it appearing in print.

Thanks to my advisor Jeff Lagarias for pointing me to the area of dualizable circle packings.

The sections that follow are not directly related to the main results of this thesis, but are here for reasons of exposition and motivation.

2.12 Euclidean polyhedra

Circle packings and Euclidean polyhedra are intimately related. The following is equivalent to the Dualizable Circle Packing Theorem 2.57:

Theorem 2.61. *Let $X = (V, E, F)$ be a polytopal decomposition of the 2-sphere $\mathbb{S}^2$. Then there exists a Euclidean polyhedron P in $\mathbb{R}^3$ which is combinatorially equivalent to X, so that every edge of P is tangent to the unit sphere. Furthermore P is uniquely determined by X, up to action by linear transformations of $\mathbb{R}^3$ fixing the unit sphere set-wise.*

To see how, let X^* be the polytopal decomposition of $\mathbb{S}^2$ dual to X. Let $\mathcal{P}^*$ be the dualizable circle packing on $\hat{\mathbb{C}}$ realizing X^*. Identify $\hat{\mathbb{C}}$ isometrically with the unit sphere in $\mathbb{R}$. Then for every disk D_f^* of $\mathcal{P}^*$, with $f \in F$, let H_f be the half-space in $\mathbb{R}^3$ whose boundary plane contains ∂D_f^*, so that the interior of the disk $D_f^* \subset \hat{\mathbb{C}}$ lies outside of H_f. Then $P = \cap_{f \in F} H_f$. Note that the group of linear transformations of $\mathbb{R}^3$ fixing the unit sphere set-wise is isomorphic to the group generated by Möbius and anti-Möbius transformations, giving us the rigidity portion of Theorem 2.61.

Theorem 2.61 contrasts sharply with the fact that there are convex polyhedra $P \subset \mathbb{R}^3$ so that there is no combinatorially equivalent polyhedron all of whose vertices lie on the unit sphere, c.f. [Ste28]. Theorem 2.61 strongly generalizes Steinitz's Theorem, which asserts only the existence of a convex polytope in $\mathbb{R}^3$ combinatorially equivalent to a given polytopal decomposition of the 2-sphere. For further discussion on this topic see [Sch87; Zie95, Chapter 4].

Like the DCPT 2.57, Theorem 2.61 has an uncertain history. For details, we refer the reader to the survey [Sac94] by Sachs, where he states both Theorem 2.61 and the DCPT 2.57, and notes their equivalence. There he writes that Theorem 2.61 was conjectured independently many times from the late 1970s through the early 1990s, giving specific references. He also speculates that Koebe himself anticipated Theorem 2.61, citing comments made in [Koe36, p. 162].

Many articles, including [BS93, Moh93, Sac94], attribute the first proof of Theorem 2.61 in its given form to Pulleyblank and Rote, and cite a manuscript in preparation. However, this manuscript seems never to have appeared in print. The special case of Theorem 2.61 where X is a triangulation of the 2-sphere is given in Thurston's lecture notes, c.f. [Thu80, Corollary 13.6.3].

Confusingly, in [Sch92], Schramm writes that Koebe proves the special case of Theorem 2.61 where X triangulates the sphere in [Koe36]. He then attributes the proof of the general case of Theorem 2.61 to Thurston, citing only the lecture notes

[Thu80], which do not contain a statement of the general case as far as we can tell.

For interest's sake, we give the following amazing generalization of Theorem 2.61, c.f. [Sch92, Theorem 1.1], proved by Schramm in his article titled, *How to cage an egg*:

Theorem 2.62. *Let X be a polytopal decomposition of the sphere, and let $E \subset \mathbb{R}^3$ be a smooth strictly convex body. Then there exists a convex polyhedron $P \subset \mathbb{R}^3$ combinatorially equivalent to X so that every edge of P is tangent to ∂E.*

His proof is not constructive. He also proves rigidity conditions on the space of such polyhedra given a fixed X, analogous to the uniqueness portion of Theorem 2.61.

2.13 Circle-packable Riemann surfaces

The Discrete Uniformization Theorem 2.15 fixes a conformal structure for a triangulation of a topological surface in a canonical way. A natural question arises in light of this observation:

Question 2.63. *Given a triangulation X of a topological surface, how can we determine which conformal structure the Discrete Uniformization Theorem 2.15 gives to X?*

Of course, there is only one conformal structure on a sphere, so Question 2.63 is trivial in this case. Let us next consider the case when X triangulates a topological open disk. In this case He and Schramm answer Question 2.63 in terms of some discrete-analytic information about the 1-skeleton of X, specifically a discretized version of conformal modulus or extremal length, see [HS95]. They also relate the circle-packing type of X to whether the random walk on the 1-skeleton of X is recurrent or transient. Beardon and Stephenson in [BS91a] show that the locally finite circle packing realizing X lies in $\mathbb{C}$ if every vertex has degree at most 6, and lies in $\mathbb{H}^2$ if every vertex had degree at least 7. These results are improved on in [HS95].

Remark 2.64. One may ask the analogous question to Question 2.63 in our setting, that is, for disk configurations realizing pairs (G, Θ) with $\Theta : E \to [0, \pi)$. When G triangulates an open disk, an answer in a similar vein is given by He in [He99, Uniformization Theorem 1.3].

Not much is known about the answer to Question 2.63 in the case where X is multiply connected.

Another natural question is:

Question 2.65. *Given a complete constant curvature Riemannian metric on an oriented topological surface S, is it possible to find a a triangulation X of S so that the Discrete Uniformization Theorem 2.15 gives X that conformal structure?*

In other words, given a pair (S, d) as in Question 2.65, we ask, is (S, d) *circle packable*? This question has been widely known since at least the mid-1980s, and remains open in the important case of compact S with positive genus. We now survey the relevant results on the subject.

There is essentially only one complete constant curvature metric d on the sphere, and it certainly gives a packable structure to the sphere. In Example 2.18 we saw that both the complex and the hyperbolic planes are packable. Thus we have answered Question 2.65 for simply connected S. In [Wil03], Williams shows that if S is open, then every pair (S, d) is packable. The idea of the proof is that the cusps of (S, d) have enough area that we may build our packing by adding disks successively, and "shuffling around" disks using the open area whenever necessary in order to fit new disks, obtaining the desired packing filling (S, d) in the limit.

The situation for compact S having positive genus is not so well understood. For a compact S, there are uncountably many d so that (S, d) have distinct conformal structures. On the other hand, there are only countably many d that give a packable pair (S, d). To see why, in light of the uniqueness part of the DUT 2.15, we need only to show that there are countably many triangulations of a compact S. This is because there are finitely many triangulations having any fixed number of faces, since there are only finitely many combinatorially distinct ways to identify the edges of those faces. In [Bro86], Brooks proves that for compact S the set of packable (S, d) is dense in the moduli space of conformal structures on S. His proof appeals to the theory of quasi-conformal deformation of Kleinian groups. The case where S is a torus is not so bad to approach with our bare hands, so we work it out as an example.

Example 2.66. If S is a torus, it is not too hard to see by elementary means that the set of circle packable (S, d) is dense in the moduli space of conformal tori. First, recall that every (S, d) is conformally equivalent to $\mathbb{C}/\langle 1, \tau \rangle$, where $\langle 1, \tau \rangle$ is the lattice generated by 1 and some τ in the upper half-plane. Furthermore, if τ_1 and τ_2 are close as complex numbers, then $\mathbb{C}/\langle 1, \tau_1 \rangle$ and $\mathbb{C}/\langle 1, \tau_2 \rangle$ are close in the moduli space of conformal tori, and conversely. Suppose that we wish to approximate the conformal torus $\mathbb{C}/\langle 1, \tau \rangle$. Recall that $\mathcal{P}_{6,\varepsilon}$ is the penny-packing of the complex plane by disks of radius ε. Restrict our attention to those ε which are of the form $1/n$ for positive integer n. Then we may normalize $\mathcal{P}_{6,\varepsilon}$ so that there are disks centered both at the

origin and at 1. Then by picking ε sufficiently small, we may ensure that there is a disk D of $\mathcal{P}_{6,\varepsilon}$ centered as close to τ as desired. Then the torus $\mathbb{C}/\langle 1, \text{center}(D)\rangle$ is packable, in particular by the image packing of $\mathcal{P}_{6,\varepsilon}$.

The study of circle-packable Riemann surfaces leads to connections from circle packing to the fields of model theory and real algebraic geometry. Specifically, we can apply the following theorem to the study of circle packing:

Tarski's Theorem 2.67. *Suppose that S is a first-order sentence in the theory of real-closed fields. If S is true in one real-closed field, then it is true in every real-closed field.*

A *real-closed field* is a field which is not algebraically closed, but which becomes algebraically closed if we adjoin a square root of -1. The only two examples we care about here are the real numbers $\mathbb{R}$ and the real algebraic numbers. A *first-order sentence in the theory of real-closed fields* is, roughly speaking, a finite logical statement, in which we may use the usual logical connectives and quantifiers ($\Rightarrow$, $\Leftrightarrow$, $\neg$, $\vee$, $\wedge$, $\forall$, $\exists$), as well as the symbols from the theory of real-closed fields (1, 0, $\times$, $+$, $=$, $>$), so that all variables in the statement are appropriately quantified. The inequality symbol $>$ is allowed because it is a theorem that every real-closed field admits a total order compatible with its field structure in the usual sense. Swan provides an expository article on Tarski's Theorem 2.67, including a proof, in [Swa05], and cites the book [KK67] of Kreisel and Krivine as his source.

We use Tarski's Theorem 2.67 to sketch an original proof of the following theorem:

Theorem 2.68. *Suppose that $\mathbb{C}/\langle 1, \tau\rangle$ is a circle-packable torus. Then τ is an algebraic number.*

The idea of the proof is to show that the existence of a doubly-periodic packing filling the complex plane, having appropriate specified combinatorics, can be expressed as an first-order sentence in the theory of real-closed fields. Given a triangulation of the plane $X = (V, E, F)$ the existence of a packing realizing X is equivalent to a statement of the following form, where the unrestricted existential quantifier ranges over $\mathbb{R}$:

$$\begin{aligned}\forall v_i \in V, \exists x_i, y_i, r_i, \big[&\forall \langle v_j, v_k\rangle \in E, (x_j - x_k)^2 + (y_j - y_k)^2 = (r_j + r_k)^2\big] \\ &\wedge \big[\forall v_j \in V : r_j > 0\big] \\ &\wedge \big[\forall v_j, v_k \in V, (v_j \neq v_k \wedge \langle v_j, v_k\rangle \notin E) \\ &\qquad \Rightarrow (x_j - x_k)^2 + (y_j - y_k)^2 > (r_j + r_k)^2\big]\end{aligned}$$

Suppose that $\mathbb{C}/\langle 1, \tau \rangle$ is packable by $\mathcal{P}$, and let X be the triangulation obtained by lifting to the plane the triangulation that $\mathcal{P}$ realizes. Then one essentially uses the periodicity of X to eliminate all but finitely many of the variables from the above logical system, arriving at a first-order sentence, call it s, in the theory of real-closed fields, which is still equivalent to the existence of the desired doubly-periodic packing. Then by the DUT 2.15 one obtains the truth of s over $\mathbb{R}$, and furthermore that there is, after a suitable normalization, a unique set of values for the x_i, y_i, r_i in $\mathbb{R}$ that satisfy s. But by Tarski's Theorem 2.67, the sentence s is also true if we allow the unrestricted existential quantifier to range only over the real algebraic numbers, so by their essential uniqueness the x_i, y_i, r_i must all be algebraic, and the claim follows.

Another proof of Theorem 2.68, which is along similar lines but uses real algebraic geometry rather than model theory, appears in [McC96, Chapter 8]. A common guess upon hearing this result is that perhaps the circle packable $\mathbb{C}/\langle 1, \tau \rangle$ are those for which τ is *constructible*, meaning that τ lies in a tower of degree 2 field extensions of $\mathbb{Q}$. However, explicit examples of non-constructible τ are not hard to find by hand, and some are given in [McC96, Chapter 8].

We remark that Tarski's Theorem 2.67 can be applied along similar lines to prove other statements about algebraicity in circle packing. We state the following theorem as an example:

Theorem 2.69. *Every dualizable circle packing on the Riemann sphere is equivalent via a Möbius transformation to one all of whose radii, tangency points, dual radii, centers, etc., are algebraic.*

Theorem 2.69 has been proved independently by Louder and Souto, using the ideas of Thurston's rigidity proof via the MPRT 2.58 together with an algebraicity result [Thu80, Proposition 6.7.4] in hyperbolic three-manifolds, also due to Thurston. At the time of writing their proof has not appeared in print.

More recently, Kojima, Mizushima, Tan, and others have studied the closely related question of which projective structures on surfaces are packable. An excellent recent survey of results and open problems in this area is [KMT06].

Finally, thanks to Chris Hall for referring us to Tarski's Theorem 2.67. Thanks to Sergiy Merenkov for referring us to [McC96]. Also thanks to Andreas Blass, Dan Hathaway, and Scott Schneider for helpful references and discussion on the model theory content of this section.

2.14 Further references

Circle packing is a broad field, and there is no hope for an introduction such as this one to be comprehensive.

- The article [Sac94] is a short survey covering parts of the area which nowadays are classical, all of which we address here to some extent. It also gives a rough outline of the history of circle packing through 1994.
- The first half of [Roh11] is an excellent survey by Rohde focused on the contributions of Oded Schramm. In particular, Schramm proved many general theorems about packing shapes other than closed disks, and Rohde provides references to the papers containing these theorems. He also gives a long list of successful applications of circle packing to other areas of math in his Section 2.2.
- Stephenson's book [Ste05] provides a more detailed, elementary, and mostly self-contained introduction to the area and could serve as a kind of "first course in circle packing."

Beyond this, whenever possible, we have given references to other more narrowly focused surveys in the text, as well as references to seminal papers on sub-fields of circle packing which were not discussed at length here.

Chapter 3

Introduction to fixed-point index

A *Jordan curve* is a homeomorphic image of a topological circle $\mathbb{S}^1$ in the complex plane $\mathbb{C}$. A *Jordan domain* is a bounded open set in $\mathbb{C}$ with Jordan curve boundary. We use the term *closed Jordan domain* or *compact Jordan domain* to refer to the closure of a Jordan domain.

Suppose that K and $\tilde{K}$ are closed Jordan domains in the plane. We define the *positive orientation* on a Jordan curve as usual. That is, as we traverse ∂K in what we call the *positive* direction, the interior of K stays to the left. From now on whenever we refer to the boundary of a Jordan domain we will suppose that it has been oriented positively.

Let K and $\tilde{K}$ be closed Jordan domains. Let $f : \partial K \to \partial \tilde{K}$ be a homeomorphism of Jordan curves which is fixed-point-free and orientation-preserving. We call such a homeomorphism *indexable*. Then $\{f(z) - z\}_{z \in \partial K}$ is a closed curve in the plane which misses the origin. It has a natural orientation induced by traversing ∂K positively. We define the *fixed-point index* of f, denoted $\eta(f)$, to be the winding number of $\{f(z) - z\}_{z \in \partial K}$ around the origin. Two examples are shown in Figures 3.1 and 3.2.

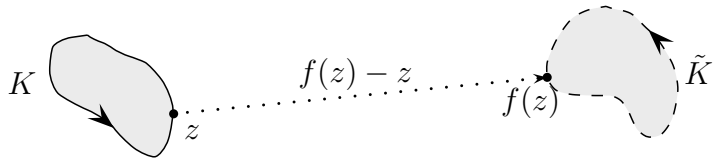

Figure 3.1: **Two closed Jordan domains K and $\tilde{K}$ so that any indexable homeomorphism $f : \partial K \to \partial \tilde{K}$ satisfies $\eta(f) = 0$.** The arrows on ∂K and $\partial \tilde{K}$ indicate the positive orientations on these Jordan curves. In this case f is indexable so long as it is orientation-preserving; the fixed-point-free condition is automatic because ∂K and $\partial \tilde{K}$ do not meet. The dashed arrow represents a vector of the form $f(z) - z$. The vector $f(z) - z$ must always point "to the right," so the curve $\{f(z) - z\}_{z \in \partial K}$ has winding number 0 around the origin, thus $\eta(f) = 0$.

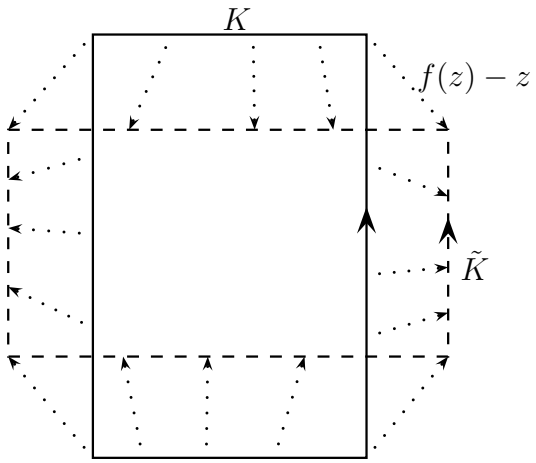

Figure 3.2: **An indexable homeomorphism $f : \partial K \to \partial \tilde{K}$ so that $\eta(f) = -1$.** Suppose we insist that f identifies the respective corners as shown. Then tracing the path of the dashed vectors $f(z) - z$ as z traverses ∂K positively, we see that $f(z) - z$ winds once clockwise around the origin, thus $\eta(f) = -1$.

We remark that the fixed-point index depends crucially on the choice of homeomorphism, and also on the way that the sets K and $\tilde{K}$ sit on top of each other. It is a worthwhile exercise to construct an indexable homeomorphism $\partial K \to \partial \tilde{K}$, for K and $\tilde{K}$ as in Figure 3.2, having fixed-point index unequal to -1.

Fixed-point index is our main technical tool. In this chapter, we prove several fundamental lemmas on fixed-point index. We then prove rigidity of circle packings which are locally finite in $\mathbb{C}$, having graphs triangulating a topological open disk, as an application. Finally, we state our main technical result, the Index Theorem 4.

3.1 The Circle Index Lemma

The following lemma can be found in [HS93, Lemma 2.2]. There it is indicated that a version of the lemma appears in [Str51]. It also appears in [Ste05, Lemma 8.13]. The moral of this lemma, and the form in which it is best remembered, is that "the fixed-point index between two circles is always non-negative."

Circle Index Lemma 3.1. *Let K and $\tilde{K}$ be closed Jordan domains in $\mathbb{C}$, with boundaries oriented positively, and let $f : \partial K \to \partial \tilde{K}$ be an indexable homeomorphism. Then the following hold.*

1. *We have $\eta(f) = \eta(f^{-1})$.*
2. *If $K \subseteq \tilde{K}$ or $\tilde{K} \subseteq K$, then $\eta(f) = 1$.*
3. *If K and $\tilde{K}$ have disjoint interiors, then $\eta(f) = 0$.*

4. If ∂K and $\partial\tilde{K}$ intersect in exactly two points, then $\eta(f) \geq 0$.

As a consequence of the above, if K and $\tilde{K}$ are closed disks in the plane, then $\eta(f) \geq 0$.

For now we give only intuitive arguments in support of Lemma 3.1, based on the proof given by He and Schramm in [HS93]. We will give a quick and easy, but totally unintuitive proof later using some machinery, c.f. Section 6.2.

"Proof." (1) By definition $\eta(f^{-1})$ is the winding number of $\{f^{-1}(\tilde{z}) - \tilde{z}\}_{\tilde{z}\in\partial\tilde{K}}$ around the origin, which is equal to the winding number of $\{z - f(z)\}_{z\in\partial K}$ around the origin under the coordinate change $f(z) = \tilde{z}$. But it is easy to see that the winding number around the origin of the curve $\{\gamma(t)\}_{t\in\mathbb{S}^1}$ is the same as the winding number around the origin of $\{-\gamma(t)\}_{t\in\mathbb{S}^1}$.

Part (2) is believable if we imagine K to be "very small," and contained in $\tilde{K}$. Then the endpoint z of the vector $f(z) - z$ does not move very much as z traverses ∂K, while the endpoint $f(z)$ of the same vector "winds once positively around K." For part (3) we refer to the example given in Figure 3.1.

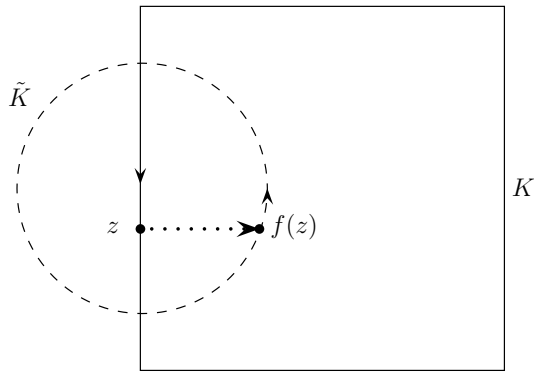

Figure 3.3: **Closed Jordan domains whose boundaries cross at exactly two points.** If $f(z) - z \in \mathbb{R}_+$ as shown, then because of the orientations on ∂K and $\partial\tilde{K}$ the endpoint z is locally moving down, and the endpoint $f(z)$ is locally moving up, so $f(z) - z$ is locally turning counter-clockwise.

For part (4) we may assume by parts (2) and (3) that K and $\tilde{K}$ are the square and circle depicted in Figure 3.3, c.f. Lemma 5.2. We consider when it is possible that the vector $f(z) - z$ points in the positive real direction, as in Figure 3.3. If $z \in \partial K$ does not lie in the interior of $\tilde{K}$, the vector $f(z) - z$ has either an imaginary component, or a negative real component. Similarly if $f(z) \in \partial\tilde{K}$ does not lie inside of K, then $f(z) - z$ has a negative real component. We conclude that the only way that $f(z) - z$ can be real and positive is if z lies along ∂K in the interior of $\tilde{K}$ and $f(z)$ lies along $\partial\tilde{K}$ inside of K. But in this case because of the orientations on ∂K and $\partial\tilde{K}$, the

vector $f(z) - z$ is locally turning in the positive direction. Thus whenever the curve $\{f(z) - z\}_{z \in \partial K}$ crosses the positive real axis it is turning in the positive direction, so this curve's total winding number around the origin cannot be negative. □

3.2 Fixed-point index additivity

The moral of this lemma is that fixed-point indices "add nicely."

Index Additivity Lemma 3.2. *Suppose that K and L are interiorwise disjoint closed Jordan domains which meet along a single positive-length Jordan arc $\partial K \cap \partial L$, similarly for $\tilde{K}$ and $\tilde{L}$. Then $K \cup L$ and $\tilde{K} \cup \tilde{L}$ are closed Jordan domains.*

Let $f : \partial K \to \partial \tilde{K}$ and $g : \partial L \to \partial \tilde{L}$ be indexable homeomorphisms. Suppose that f and g agree on $\partial K \cap \partial L$. Let $h : \partial(K \cup L) \to \partial(\tilde{K} \cup \tilde{L})$ be induced via restriction to f or g as necessary. Then h is an indexable homeomorphism and $\eta(h) = \eta(f) + \eta(g)$.

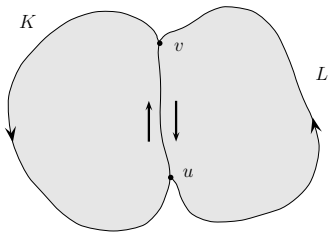

Figure 3.4: **An illustration of fixed-point index additivity.**

Proof. The situation is as depicted in Figure 3.4. We may consider $\eta(f)$ to be $1/2\pi$ times the change in argument of the vector $f(z) - z$, as z traverses ∂K once in the positive direction. Then as z varies positively in ∂K and in $\partial \tilde{K}$ the contributions to the sum $\eta(f) + \eta(g)$ along $\partial K \cap \partial L$ cancel. □

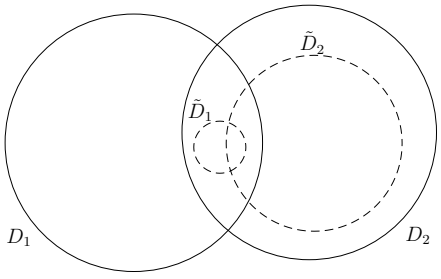

Figure 3.5: **A computational example applying the Circle Index Lemma 3.1 and fixed-point index additivity.** Suppose $f_1 : \partial D_1 \to \partial \tilde{D}_1$ and $f_2 : \partial D_2 \to \partial \tilde{D}_2$ are indexable homeomorphisms that agree on $\partial D_1 \cap \partial D_2$. Then f_1 and f_2 induce indexable homeomorphisms $g : \partial(D_2 \setminus D_1) \to \partial(\tilde{D}_2 \setminus \tilde{D}_1)$ and $h : \partial(D_1 \cup D_2) \to \partial(\tilde{D}_1 \cup \tilde{D}_2)$. Part (2) of the Circle Index Lemma 3.1 gives that $\eta(h) = \eta(f_1) = 1$, so by index additivity $\eta(g) = 0$.

3.3 Stability of fixed-point index under small perturbations

The following lemma says essentially that given $f : \partial K \to \partial\tilde{K}$ an indexable homeomorphism, we may move K around a "little bit," leaving $\tilde{K}$ fixed, without affecting the fixed-point index $\eta(f)$. This lemma appears in [Ste05, Lemma 8.11].

Index Stability Lemma 3.3. *Let K and $\tilde{K}$ be closed Jordan domains and let $f : \partial K \to \partial\tilde{K}$ be an indexable homeomorphism. Then there is an $\varepsilon > 0$ so that the following holds. Let $\phi : \partial K \to \mathbb{C}$ be an orientation-preserving homeomorphism onto its image. Suppose further that ϕ satisfies $|\phi(z) - z| < \varepsilon$ for all $z \in \partial K$. Then $f^* : \phi(\partial K) \xrightarrow{\phi^{-1}} \partial K \xrightarrow{f} \partial\tilde{K}$ is an indexable homeomorphism, and $\eta(f) = \eta(f^*)$.*

Proof. Let $\varepsilon = \inf_{z\in\partial K} |f(z) - z|$. Then $\varepsilon > 0$ because ∂K is compact and the function $|f(z) - z|$ is continuous and non-zero for $z \in \partial K$. For $z \in \partial K$, let $\phi_t(z) = tz + (1-t)\phi(z)$, so that ϕ_1 is the identity on ∂K and $\phi_0 = \phi$. Let γ_t be the closed curve $\{f(z) - \phi_t(z)\}_{z\in\partial K}$. Now $|f(z) - \phi_t(z)| \geq |f(z) - z| - |z - \phi_t(z)|$ by the triangle inequality, but $|f(z) - z| \geq \varepsilon > |z - \phi(z)| \geq |z - \phi_t(z)|$, so $|f(z) - \phi_t(z)| > 0$ for all $z \in \partial K$ and $t \in [0, 1]$. Thus every curve γ_t misses the origin. Furthermore the curves γ_t vary continuously in t, so all of them have the same winding number around the origin. But the winding number of γ_0 around the origin is $\eta(f^*)$, and the winding number of γ_1 around the origin is $\eta(f)$. □

Let K and $\tilde{K}$ be closed Jordan domains. We say that K and $\tilde{K}$ are in *general position* if ∂K and $\partial\tilde{K}$ "cross" wherever they meet. More precisely, we say that K and $\tilde{K}$ are in *general position* for any $z \in \partial K \cap \partial\tilde{K}$, locally near z the curves ∂K and $\partial\tilde{K}$ are topologically as in Figure 3.6a. That is, there is an open neighborhood U of z and a homeomorphism $\phi : U \to \mathbb{D}$ sending $\partial K \cap U$ to $\mathbb{R} \cap \mathbb{D}$ and sending $\partial\tilde{K} \cap U$ to $i\mathbb{R} \cap \mathbb{D}$. For example, the general position hypothesis on K and ∂K precludes intersection points $\partial K \cap \partial\tilde{K}$ which locally look like Figures 3.6b and 3.6c. We remark that in general, Jordan curves may meet in other local configurations than the three given in Figure 3.6.

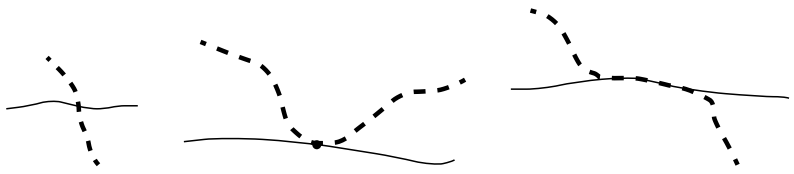

Figure 3.6: **Various ways that two Jordan curves ∂K and $\partial\tilde{K}$ can meet.** The solid arcs represent pieces of ∂K and the dashed ones pieces of $\partial\tilde{K}$.

Remark 3.4. Actually, in general, Jordan curves are poorly behaved, and a "random" pair of Jordan curves, if they meet, will probably not be in what we call general position. However, all examples of Jordan curves that we ever consider are piecewise smooth, and a "random" pair of piecewise smooth Jordan curves probably will be in general position, so our terminology makes sense in this context. Thanks to Steffen Rohde for pointing out this issue.

3.4 Moving three points

The following lemma says essentially that we may almost always prescribe the images of three points on ∂K in $\partial \tilde{K}$, and obtain an indexable homeomorphism $\partial K \to \partial \tilde{K}$ with non-negative index, which respects this prescription.

Three Point Prescription Lemma 3.5. *Let K and $\tilde{K}$ be compact Jordan domains in general position. Let $z_1, z_2, z_3 \in \partial K \setminus \partial \tilde{K}$ appear in counterclockwise order, similarly $\tilde{z}_1, \tilde{z}_2, \tilde{z}_3 \in \partial \tilde{K} \setminus \partial K$. Then there is an indexable homeomorphism $f : \partial K \to \partial \tilde{K}$ sending $z_i \mapsto \tilde{z}_i$ for $i = 1, 2, 3$, so that $\eta(f) \geq 0$.*

A version of this lemma is stated in [Ste05, Lemma 8.14], but the proof given there is incomplete. However, that proof works in "almost all" cases, and anyway is a quick heuristic for the truth of the TPPL 3.5, so we sketch it here. It depends on the following well-known theorem (see [Car13] for the original paper):

Carathéodory's Theorem 3.6. *If Ω and $\tilde{\Omega}$ are open Jordan domains in the plane, then any biholomorphism $\varphi : \Omega \to \tilde{\Omega}$ extends uniquely to a homeomorphism $\bar{\varphi} :$ closure$(\Omega) \to$ closure$(\tilde{\Omega})$. Furthermore if $z_1, z_2, z_3 \in \partial\Omega$ are distinct, and $\tilde{z}_1, \tilde{z}_2, \tilde{z}_3 \in \partial\tilde{\Omega}$ are distinct, and both triples occur in the same order around their boundaries, then φ may be chosen uniquely so that $\bar{\varphi}$ sends z_1, z_2, z_3 to $\tilde{z}_1, \tilde{z}_2, \tilde{z}_3$ respectively.*

Then let φ be the Riemann mapping from the interior of K to the interior of $\tilde{K}$, c.f. the Riemann Mapping Theorem 2.4. Then by Carathéodory's Theorem 3.6 the map φ limits uniquely to a boundary homeomorphism $f : \partial K \to \partial \tilde{K}$. Choose φ so that f sends z_1, z_2, z_3 to $\tilde{z}_1, \tilde{z}_2, \tilde{z}_3$ respectively.

Suppose that f does not have any fixed points. Suppose also for simplicity that ∂K and $\partial \tilde{K}$ are smooth. Then, using the complex analysis definition of winding

number (see [Ahl78, Section 2.1; Rud87, p. 204]), we have that:

$$\begin{aligned}\eta(f) &= \oint_{\{f(z)-z\}_{z\in\partial K}} \frac{dw}{w} \\ &= \oint_{\partial K} \frac{\varphi'(w)-1}{\varphi(w)-w} dw\end{aligned}$$

Then by the Argument Principle (see [Ahl78, Theorem 5.18]) the second integral counts the number of zeros minus the number of poles of $f(z) - z$ in the interior of K, but $f(z) - z$ is holomorphic in the interior of K, so this integral is non-negative.

Actually φ' is undefined on ∂K, because φ is not holomorphic in a neighborhood of K as technically required by the Argument Principle, so the second integral does not quite make sense. It is possible to get around this problem using homotopies and the Index Stability Lemma 3.3. However, there is a more serious issue, namely that *a priori* f may have many fixed points, and it is not clear how to get rid of them, especially if we do not wish to move K and $\tilde{K}$. This is the gap in the proof given in [Ste05]. In particular, moving K and $\tilde{K}$ would require adjusting the given statement of the TPPL 3.5.

We give an original elementary inductive proof of Lemma 3.5, using only plane topology arguments, in Section 6.5.

3.5 Proof of rigidity and uniformization of circle packings

In this section we prove the following theorem, which is a special case of our Main Rigidity Theorem 1:

Theorem 3.7. *Suppose that $\mathcal{P}$ and $\tilde{\mathcal{P}}$ are circle packings in $\hat{\mathbb{C}}$, sharing a contact graph that triangulates the 2-sphere $\mathbb{S}^2$. Then $\mathcal{P}$ and $\tilde{\mathcal{P}}$ differ by a Möbius or anti-Möbius transformation.*

We include the proof as an example of the surprising power of the fixed-point index, and because it is an excellent warm-up to the proof of our main result. We then easily adapt the proof of Theorem 3.7 to obtain the following two theorems:

Theorem 3.8. *Suppose that $\mathcal{P}$ and $\tilde{\mathcal{P}}$ are circle packings locally finite in $\mathbb{C}$, sharing a contact graph that triangulates a topological open disk. Then $\mathcal{P}$ and $\tilde{\mathcal{P}}$ differ by a Euclidean similarity.*

Theorem 3.9. *There cannot be two circle packings $\mathcal{P}$ and $\tilde{\mathcal{P}}$ sharing a contact graph triangulating a topological open disk, so that one is locally finite in $\mathbb{C}$ and the other is locally finite in $\mathbb{H}^2$.*

3.5.1 Proof of Theorem 3.7

The argument is by contradiction. The main idea is to superimpose the packings $\mathcal{P}$ and $\tilde{\mathcal{P}}$ on the Riemann sphere in a convenient way. In particular, we will isolate two topological quadrilaterals Q and $\tilde{Q}$ so that any indexable homeomorphism $\partial Q \to \partial \tilde{Q}$ identifying corresponding corners has a negative fixed-point index. Then, using the Circle Index Lemma 3.1, the Index Additivity Lemma 3.2, and the Three Point Prescription Lemma 3.5, we construct an indexable homeomorphism $\partial Q \to \partial \tilde{Q}$ which identifies corresponding corners, but which has a non-negative fixed-point index by construction. This gives us our desired contradiction.

Let $X = (V, E, F)$ be the triangulation of $\mathbb{S}^2$ which $\mathcal{P}$ and $\tilde{\mathcal{P}}$ realize. Suppose for contradiction that $\mathcal{P}$ and $\tilde{\mathcal{P}}$ are not equivalent under any Möbius or anti-Möbius transformation.

For the first part of the proof, we apply a sequence of normalizations to $\mathcal{P}$ and to $\tilde{\mathcal{P}}$. Let $f_0 = \langle v_1, v_2, v_3 \rangle \in F$ be a face of X. We first normalize via a Möbius transformation so that $D_i = \tilde{D}_i$ for $i = 1, 2, 3$. Here $D_i \in \mathcal{P}$ is the disk corresponding to $v_i \in V$ as usual, similarly $\tilde{D}_i \in \tilde{\mathcal{P}}$.

Our next normalization is in our initial choice of f_0 and our labeling of the v_i, as per the following observation:

Observation 3.10. *Let v_4 denote the vertex of X other than v_1 so that $\langle v_2, v_3, v_4 \rangle$ is a face of X. Then there is some choice of $f_0 = \langle v_1, v_2, v_3 \rangle$ so that the disks D_4 and $\tilde{D}_4$ are not equal after our normalization identifying $D_i = \tilde{D}_i$ for $i = 1, 2, 3$.*

If there were no such choice of f_0 then in fact every pair of corresponding disks D_i and $\tilde{D}_i$ would coincide after our first normalization, and so $\mathcal{P}$ and $\tilde{\mathcal{P}}$ coincide.

Next, we insist that ∞ lies in the interstice formed by $D_1 = \tilde{D}_1, D_2 = \tilde{D}_2, D_3 = \tilde{D}_3$. Finally, we insist that $D_1 = \tilde{D}_1, D_2 = \tilde{D}_2, D_3 = \tilde{D}_3$ all have Euclidean radius 1, and that D_2 and D_3 are tangent at a point lying on the horizontal axis, so that D_1 lies to their left. The situation for $\mathcal{P}$ is depicted in Figure 3.7.

From now on we work in the plane $\mathbb{C}$, in the sense that $\infty \in \hat{\mathbb{C}}$ will not move again for the remainder of the proof. Note that every face f of F corresponds to some interstices $U_f \subset \mathbb{C}$ and $\tilde{U}_f \subset \mathbb{C}$ of $\mathcal{P}$ and $\tilde{\mathcal{P}}$ respectively, except for f_0, for which the interstices $U_{f_0} = \tilde{U}_{f_0}$ contain ∞. Let Q be the topological quadrilateral shown in Figure 3.7b. More precisely, let $V_Q = V \setminus \{v_1, v_2, v_3\}$, and let $F_Q = F \setminus \{f_0 = \langle v_1, v_2, v_3 \rangle, \langle v_2, v_3, v_4 \rangle\}$. Then we define $Q = \bigcup_{v \in V_Q} D_v \cup \bigcup_{f \in F_Q} U_f$. Define the analogous objects for $\tilde{\mathcal{P}}$ in the obvious way.

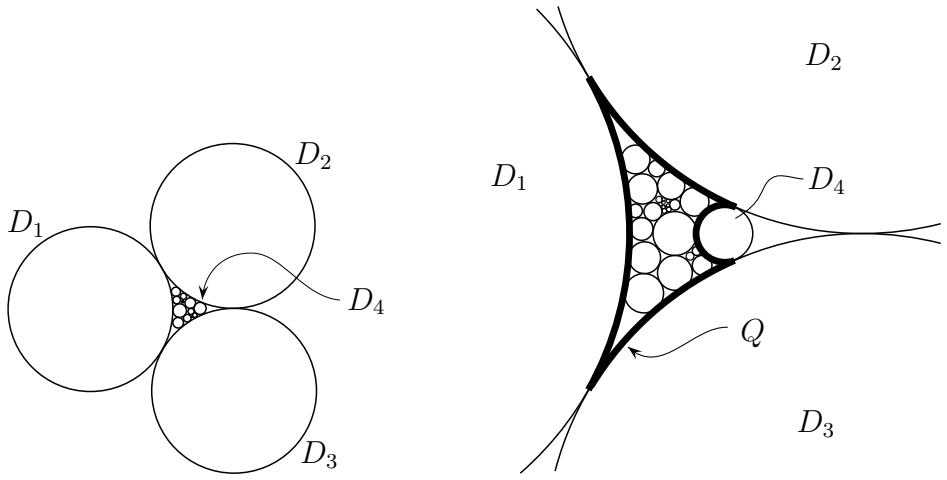

Figure 3.7: **The packing $\mathcal{P}$ after some normalizations.** The disks of $\mathcal{P}$ all lie between D_1, D_2, D_3. Note that the interstice formed by D_1, D_2, D_3 on $\hat{\mathbb{C}}$ is the outside region in these figures. The disk D_4 is "the first disk of $\mathcal{P} \setminus \{D_1, D_2, D_3\}$ we get to if we start scanning from the right." In (b) the topological quadrilateral Q is outlined in bold.

We now apply one final transformation to $\mathcal{P}$. First, suppose without loss of generality that the Euclidean radius of D_4 is larger than that of $\tilde{D}_4$. Then translate every disk of $\mathcal{P}$ to the right by a small amount $\varepsilon > 0$, leaving the disks of $\tilde{\mathcal{P}}$ unchanged. Denote this transformation by T_ε. We will discuss more precise requirements on ε later. The situation is depicted in Figure 3.8. The essential point is that there is an open interval of values that $\varepsilon > 0$ may take so that after all of our transformations, the topological quadrilaterals Q and $\tilde{Q}$ are arranged qualitatively as in Figure 3.8b. In particular, in this case we get that any indexable homeomorphism $\partial Q \to \partial \tilde{Q}$ identifying corresponding corners has a negative fixed-point index.

We now move on to the second part of the proof. We hope to obtain an indexable homeomorphism $\varphi_Q : \partial Q \to \partial \tilde{Q}$ which identifies corresponding corners, constructed in such a way that we can prove that $\eta(\varphi_Q) \geq 0$. First, for any $f = \langle u_1, u_2, u_3 \rangle \in F_Q$, a homeomorphism $\varphi_f : \partial U_f \to \partial \tilde{U}_f$ is called *faithful* if it identifies corresponding corners, equivalently if it restricts to homeomorphisms $\partial U_f \cap \partial D_{u_i} \to \partial \tilde{U}_f \cap \partial \tilde{D}_{u_i}$ for $i = 1, 2, 3$. Go back and ensure that we chose ε so that every pair of corresponding interstices is in general position. We may do so because there are only countably many choices of ε for which this fails, and there were uncountably many acceptable choices for ε until this point, so there continue to be uncountably many acceptable choices after we impose this extra condition. Then for all $f \in F \setminus \{\langle f_0 = v_1, v_2, v_3 \rangle\}$,

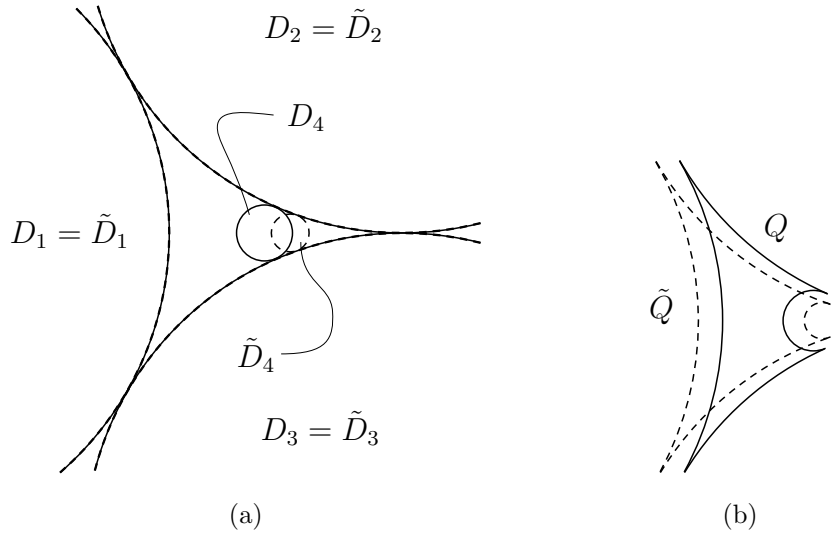

Figure 3.8: **The interaction between $\mathcal{P}$ and $\tilde{\mathcal{P}}$ before and after applying T_ε.** In (a) we see the superimposition of the D_i with the $\tilde{D}_i$ before applying T_ε to $\mathcal{P}$. The disks D_i are drawn solid, and the disks $\tilde{D}_i$ are drawn dashed. In (b) we see the relative positions of Q and $\tilde{Q}$ after applying T_ε to $\mathcal{P}$.

fix a faithful indexable homeomorphism $\varphi_f : \partial U_f \to \partial \tilde{U}_f$ so that $\eta(\varphi_f) \geq 0$, via the Three Point Prescription Lemma 3.5. Then:

Observation 3.11. *For every $v \in V \setminus \{v_1, v_2, v_3\}$ the homeomorphisms φ_f induce an indexable homeomorphism $\varphi_v : D_v \to \tilde{D}_v$. The φ_f also induce an indexable homeomorphism $\varphi_Q : \partial Q \to \partial \tilde{Q}$, which identifies corresponding corners.*

But then by the Index Additivity Lemma 3.2, we get:

$$\eta(\varphi_Q) = \sum_{v \in V_Q} \eta(\varphi_v) + \sum_{f \in F_Q} \eta(\varphi_f)$$

We have already observed that the left side of the sum is negative. The right side of the sum is positive, because $\eta(\varphi_v) \geq 0$ for all v by the Circle Index Lemma 3.1, so we obtain our desired contradiction. □

3.5.2 Proof of Theorem 3.8

The proof of Theorem 3.8 proceeds along the same lines, except that after our first round of normalizations identifying D_i and $\tilde{D}_i$ for $i = 1, 2, 3$, and sending a point of their interstice to ∞, the remaining disks of $\mathcal{P}$ accumulate around a point z_∞, as do those of $\tilde{\mathcal{P}}$ around a point $\tilde{z}_\infty$. The points z_∞ and $\tilde{z}_\infty$ may coincide or may be different. We define and apply T_ε as before, this time making sure that z_∞ and $\tilde{z}_\infty$

are different after applying T_ε.

Next, pick disjoint neighborhoods W and $\tilde{W}$ of z_∞ and $\tilde{z}_\infty$ respectively, and contained in Q and $\tilde{Q}$ respectively. Then, let V_L be the set of $v \in V$ so that both $D_v \subset W$ and $\tilde{D}_v \subset \tilde{W}$. Similarly let F_L be the set of $f \in F$ so that both $U_f \subset W$ and $\tilde{U}_f \subset \tilde{W}$. Let L (which stands for leftovers) be the union $\bigcup_{v \in V_L} D_v \cup \bigcup_{f \in F_L} U_f \cup z_\infty$, and define $\tilde{L}$ similarly. Then the φ_f and φ_v induce an indexable homeomorphism $\varphi_L : \partial L \to \partial \tilde{L}$, and $\eta(\varphi_L) = 0$ because L and $\tilde{L}$ are disjoint. Then by the Index Additivity Lemma 3.2, we get:

$$-1 = \eta(\varphi_Q) = \eta(\varphi_L) + \sum_{v \in V_Q \setminus V_L} \eta(\varphi_v) + \sum_{f \in F_Q \setminus F_L} \eta(\varphi_f) \geq 0$$

Thus we have our desired contradiction. □

3.5.3 Proof of Theorem 3.9

The adaptation here is along similar lines as the adaptation to prove Theorem 3.8. Let $\mathcal{P}$ and $\tilde{\mathcal{P}}$ be as in the statement of the theorem. Suppose for contradiction that $\mathcal{P}$ is locally finite in $\mathbb{C}$, and $\tilde{\mathcal{P}}$ is locally finite in $\mathbb{H}^2 \cong \mathbb{D}$. This time, after our normalizations, the disks of $\mathcal{P}$ accumulate around a single point z_∞, and the disks of $\tilde{\mathcal{P}}$ accumulate around some circle C contained in the bounded region in the plane formed between $D_1 = \tilde{D}_1, D_2 = \tilde{D}_1, D_3 = \tilde{D}_3$. This time, ensure that we chose ε so that z_∞ does not lie on the circle C. For $\delta > 0$ let W_δ be the neighborhood of radius δ around z_∞, and let $\tilde{W}_\delta$ be the set consisting of all points within distance δ of a point on the circle C. Choose δ small enough that W_δ and $\tilde{W}_\delta$ are disjoint. Let V_L and F_L be defined as before. Let D be the closed disk in $\mathbb{C}$ bounded by C. By deleting finitely many vertices from V_L and finitely many faces from F_L if necessary, we may suppose that L, defined as before, is a closed Jordan domain, as is $\tilde{L} := \bigcup_{v \in V_L} \tilde{D}_v \cup \bigcup_{f \in F_L} \tilde{U}_f \cup D$. Also ∂L and $\partial \tilde{L}$ are disjoint, so $\eta(\varphi_L) \geq 0$ for any indexable $\varphi_L : \partial L \to \partial \tilde{L}$ by the Circle Index Lemma 3.1. The proof finishes as before. □

Remark 3.12. At first glance, it may seem like a similar adaptation can be used to prove rigidity of packings which are locally finite in $\mathbb{H}^2$ and have contact graphs triangulating a topological open disk. However, the situation is subtle. The idea would be to use the Circle Index Lemma 3.1 on the circles which are the images under normalization of the boundaries of the hyperbolic spaces where the two packings live. However, it is not clear (although it is true) that the two circle packings induce a

homeomorphism of the boundaries of their supporting hyperbolic spaces in the natural way. It is also unclear how one would get rid of any fixed points which may arise in the construction of φ_L in such a proof.

3.6 The Incompatibility Theorem

The essential topological idea of the proofs of this part of the thesis originates with the so-called Incompatibility Theorem of Oded Schramm. It says roughly the following. We will pack topological rectangles with shapes in such a way that every region left over is a curvilinear triangle. Two examples of such packings are given in Figure 3.9. Suppose, as in the figure, that we pack two topological rectangles R and $\tilde{R}$ this way, so that the contact graphs of the shapes are the same, and so that given two corresponding corners z of R and $\tilde{z}$ of $\tilde{R}$, the shape K_z in R nearest z and the shape $\tilde{K}_{\tilde{z}}$ in $\tilde{R}$ nearest $\tilde{z}$ correspond to one another in the identification of the contact graphs of the packings. In other words, the packings agree about their contacts and about which shape is in which corner. Suppose that we overlay the rectangles R and $\tilde{R}$ so that they cut each other. Then there is a pair of corresponding shapes which cut each other. Two examples are shown in Figure 3.10. Here we say that two sets A and B *cut each other* if $A \setminus B$ or $B \setminus A$ is disconnected.

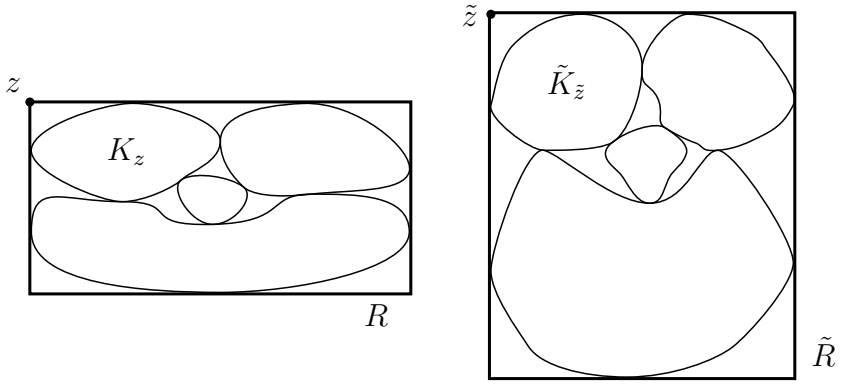

Figure 3.9: **Two topological rectangles packed with shapes** so that every remaining region is a curvilinear triangle.

Theorems 3.7, 3.8, and 3.9 can be proved via the Incompatibility Theorem by applying the same normalizations and transformations as in our proof, and then observing that two circles can never cut each other, giving us the desired contradiction. The end of our proof of Theorem 3.7 suggests a strategy for proving the Incompatibility Theorem, and indeed a careful proof can be carried out along those lines using the fixed-point index tools we have in hand so far. The difficulties that manifest

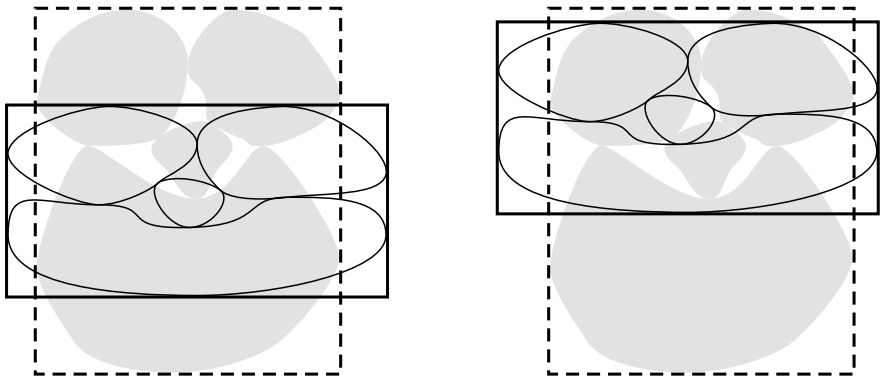

Figure 3.10: **Shapes cutting each other.** We have drawn R and $\tilde{R}$ on top of each other in two different ways, so that R and $\tilde{R}$ cut each other, and in both cases some pair of corresponding shapes of our packings cut each other.

themselves when we try to generalize this method of proof to disk configurations with overlaps will become clearer in Section 3.9, where we discuss obstructions to possible generalizations of our main technical result.

The original proof by Schramm of the Incompatibility Theorem is via different methods, is elementary, and is somewhat messier. The precise technical statement of the Incompatibility Theorem originally given by Schramm is also rather technical, which is why we do not state it here, giving instead a picture of the main idea. The original statement and proof can be found in [Sch91, Theorem 3.1], where it is used to give the first proof of Theorem 3.8.

3.7 A more general definition of fixed-point index

In this section we give the sufficiently general definition of fixed-point index required to accommodate the statement of our Main Index Theorem 4. See Figure 3.11 for a motivating example.

Definition 3.13. Suppose that K is a compact set in $\mathbb{C}$ that can be written $K = \cup_{i=1}^{n} K_i \setminus \cup_{j=1}^{m} U_j$, where

- the sets K_i are pairwise interiorwise disjoint closed Jordan domains, so that the boundaries of any two of the K_i meet at at most finitely many points, and
- the sets U_j are pairwise disjoint open Jordan domains, so that every U_j is contained in some K_i.

Suppose that the compact set $\tilde{K} = \cup_{i=1}^{n} \tilde{K}_i \setminus \cup_{j=1}^{m} \tilde{U}_j$ is such that $\tilde{K}_i$ and $\tilde{U}_j$ satisfy the same conditions. A homeomorphism $f : \partial K \to \partial \tilde{K}$ is called *indexable* if it is fixed-

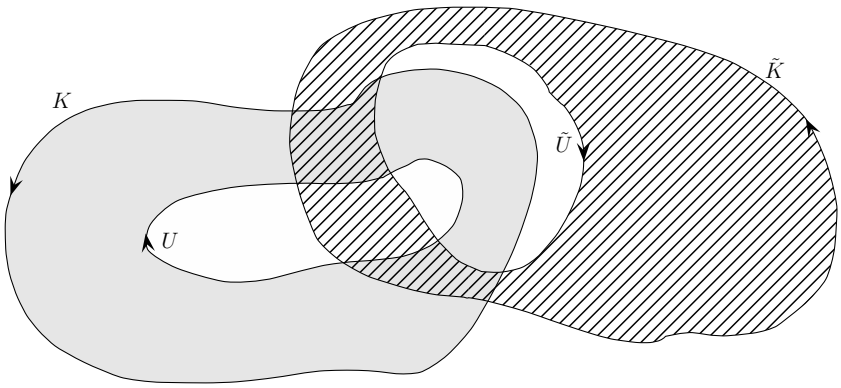

Figure 3.11: **Generalizing the definition of fixed-point index.** Denote $K = K_0 \backslash U$ where K_0 is a closed Jordan domain and U is an open Jordan domain, similarly $\tilde{K} = \tilde{K}_0 \setminus \tilde{U}$. We have shown the *positive orientation* on ∂K, that is, the orientation so that the interior of K stays to the left as ∂K is traversed *positively*. This restricts to the negative orientation on ∂U, similarly for for $\tilde{K}$ and $\tilde{U}$. Thus if $f : \partial K \to \partial \tilde{K}$ is a fixed-point-free orientation-preserving homeomorphism mapping $\partial K_0 \to \partial \tilde{K}_0$ and $\partial U \to \partial \tilde{U}$, it is natural to define $\eta(f) = \eta(\partial K_0 \xrightarrow{f} \partial \tilde{K}_0) - \eta(\partial U \xrightarrow{f} \partial \tilde{U})$.

point-free, orientation-preserving, and restricts to homeomorphisms $K_i \cap \partial K \to \tilde{K}_i \cap \partial \tilde{K}$ and $U_j \cap \partial K \to \tilde{U}_j \cap \partial \tilde{K}$ for all i, j. Then f induces indexable homeomorphisms $\partial K_i \to \partial \tilde{K}_i$ and $\partial U_j \to \partial \tilde{U}_j$ for all i, j. We define the *fixed-point index* of f to be $\eta(f) = \sum_{i=1}^{n} \eta(\partial K_i \xrightarrow{f} \partial \tilde{K}_i) - \sum_{j=1}^{m} \eta(\partial U_j \xrightarrow{f} \partial \tilde{U}_j)$.

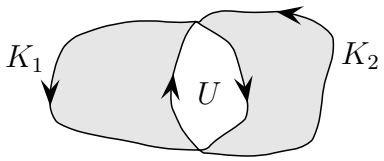

Figure 3.12: **A compact set that breaks up in two different ways, both of which are accommodated by Definition 3.13.** The shaded set can be written either as the union of two interiorwise disjoint closed Jordan domains K_1 and K_2, with no open Jordan domains removed, or as a single closed Jordan domain with an open Jordan domain U removed.

Definition 3.13 is general enough to accommodate any situation that could arise in the statement of Theorem 4. Then Theorem 4 holds using this definition of fixed-point index. We remark that a given compact set K may decompose in more than one way that satisfies the requirements of this definition, c.f. Figure 3.12. We leave it to the reader to check that the same value for the fixed-point index is obtained regardless of which decomposition is chosen. We can now also discuss fixed-point index additivity in more complicated situations, c.f. Figure 3.13. In general, fixed-

point index additivity continues to hold using Definition 3.13. We leave the details of working this out to the reader.

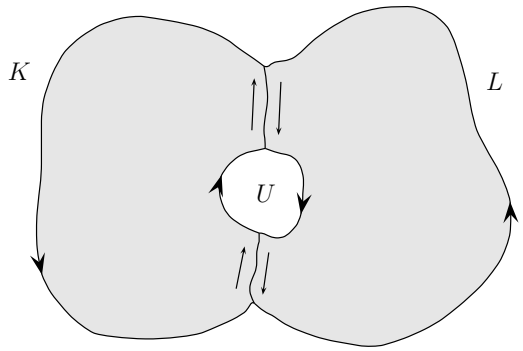

Figure 3.13: **An example of fixed-point index additivity where $K \cup L$ is not a closed Jordan domain.** Let $f : \partial K \to \partial \tilde{K}$ and $g : \partial L \to \partial \tilde{L}$ be indexable homeomorphisms which agree on $\partial K \cap \partial L$. Then the contributions from $\eta(f)$ and $\eta(g)$ along $\partial K \cap \partial L$ cancel in the sum $\eta(f) + \eta(g)$, as in the proof of Lemma 3.2. Furthermore we see that the positive orientations on ∂K and ∂L induce the negative orientation on ∂U, so the contribution to $\eta(f) + \eta(g)$ along ∂U is $-\eta(\partial U \stackrel{f,g}{\to} \partial \tilde{U})$ as desired, where $\partial U \stackrel{f,g}{\to} \partial \tilde{U}$ denotes the homeomorphism induced by restriction to f or g as necessary.

3.8 Statement of the main technical result

Let $\mathcal{K} = \{K_1, \ldots, K_n\}$ and $\tilde{\mathcal{K}} = \{\tilde{K}_1, \ldots, \tilde{K}_n\}$ be collections of closed Jordan domains. We denote $\partial\mathcal{K} = \partial \cup_{i=1}^n K_i$, similarly $\partial\tilde{\mathcal{K}} = \cup_{i=1}^n \tilde{K}_i$. A homeomorphism $f : \partial\mathcal{K} \to \partial\tilde{\mathcal{K}}$ is called *faithful* if whenever we restrict f to $K_j \cap \partial\mathcal{K}$ we get a homeomorphism $K_j \cap \partial\mathcal{K} \to \tilde{K}_j \cap \partial\tilde{\mathcal{K}}$. The particular choice of indices of K_i and $\tilde{K}_i$ is important in determining whether a given homeomorphism is faithful, so we consider the labeling to be part of the information of the collections. Note that in general $\partial\mathcal{K}$ and $\partial\tilde{\mathcal{K}}$ need not be homeomorphic, and even if they are homeomorphic there may still be no faithful homeomorphism between them.

We now give a weak form of our main technical result which is much easier to state, because it gives us a sense of where we are going, indicating roughly how we are generalizing the Circle Index Lemma 3.1.

Main Index Theorem (weak form). *Let $\mathcal{D} = \{D_1, \ldots, D_n\}$ and $\tilde{\mathcal{D}} = \{\tilde{D}_1, \ldots, \tilde{D}_n\}$ be thin disk configurations in the plane $\mathbb{C}$ sharing a contact graph $G = (V, E)$. Suppose that $\Theta : E \to [0, \pi)$, and that both $\mathcal{D}$ and $\tilde{\mathcal{D}}$ realize (G, Θ). Let $f : \partial\mathcal{D} \to \partial\tilde{\mathcal{D}}$ be a faithful indexable homeomorphism. Then $\eta(f) \geq 0$.*

This follows from the full statement of the Main Index Theorem 4 and Lemma 3.3. It can be thought of as a "first approximation" to the Index Theorem 4, in the same way that the statement that $\eta(f) \geq 0$ for indexable $f : \partial D \to \partial \tilde{D}$ where D and $\tilde{D}$ are closed disks is a "first approximation" to the Circle Index Lemma 3.1.

Moving on, we say that $\mathcal{D}$ and $\tilde{\mathcal{D}}$ are in *general position* if for all i, j we have

- that $\partial D_i \cap \partial D_j$ is disjoint from $\cup_{k=1}^{n} \partial \tilde{D}_k$,
- that $\partial \tilde{D}_i \cap \partial \tilde{D}_j$ is disjoint from $\cup_{k=1}^{n} \partial D_k$, and
- that ∂D_i and $\partial \tilde{D}_j$ are not tangent.

Next, a subset $I \subset \{1, \ldots, n\}$ is called *subsumptive* if

- either $D_i \subset \tilde{D}_i$ for every $i \in I$, or $\tilde{D}_i \subset D_i$ for every $i \in I$, and
- the set $\cup_{i \in I} D_i$ is connected, equivalently the set $\cup_{i \in I} \tilde{D}_i$ is connected.

The subsumptive set I is called *isolated* if there is no $i \in I$ and $j \in \{1, \ldots, n\} \setminus I$ so that one of $D_i \cap D_j = E_{ij}$ and $\tilde{D}_i \cap \tilde{D}_j = \tilde{E}_{ij}$ contains the other. The main technical result of this thesis is the following theorem.

Main Index Theorem 4. *Let $\mathcal{D} = \{D_1, \ldots, D_n\}$ and $\tilde{\mathcal{D}} = \{\tilde{D}_1, \ldots, \tilde{D}_n\}$ be thin disk configurations in the plane $\mathbb{C}$ in general position, sharing a contact graph $G = (V, E)$. Suppose that $\Theta : E \to [0, \pi)$, and that both $\mathcal{D}$ and $\tilde{\mathcal{D}}$ realize (G, Θ). Let $f : \partial \mathcal{D} \to \partial \tilde{\mathcal{D}}$ be a faithful indexable homeomorphism. Then $\eta(f)$ is at least the number of maximal isolated subsumptive subsets of $\{1, \ldots, n\}$. In particular $\eta(f) \geq 0$.*

For an example, look ahead to Figure 4.5 on p. 88. There we know that $\eta(f) \geq 1$ for f satisfying the hypotheses of Theorem 4.

3.9 Possible generalizations of the Index Theorem 4

First, one may hope to weaken the condition that the D_i and $\tilde{D}_i$ are closed metric disks in $\mathbb{C}$. However, Figure 3.2 shows that Theorem 4 already fails for $n = 1$ if we allow shapes whose boundaries meet at more than two points. For another example, if we allow the D_i and $\tilde{D}_i$ to be ellipses and insist on the extra condition that ∂D_i and $\partial \tilde{D}_i$ may meet at at most two points, then Figure 3.14 provides a counterexample for $n = 2$.

Next, one may hope to eliminate the thinness condition. We use this hypothesis rather strongly in our proof of Theorem 4, but are unaware of any counterexamples

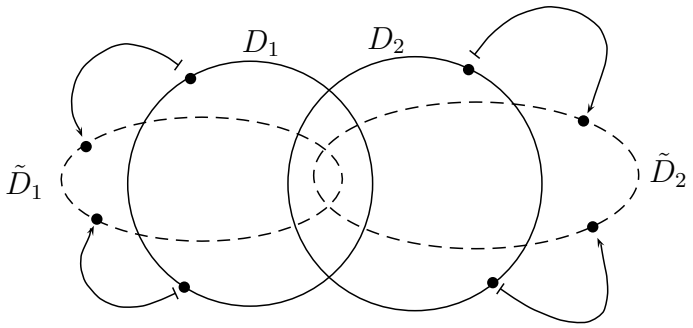

Figure 3.14: **A counterexample to Theorem 4 if we allow the $\tilde{D}_i$ to be ellipses.** We may achieve such a configuration even if we insist that $\measuredangle(\tilde{D}_1, \tilde{D}_2) = \measuredangle(D_1, D_2)$. Then any indexable $f : \partial(D_1 \cup D_2) \to \partial(\tilde{D}_1 \cup \tilde{D}_2)$ making the shown identifications gives $\eta(f) = -1$.

if it is omitted and suspect that at least the weak form of the Index Theorem 4 holds with the thinness requirement removed.

Finally, one may hope to eliminate the angle compatibility condition, but Figure 3.15 provides a counterexample. On the other hand, we do not use angle compatibility in our proof of Proposition 6.7, so Theorem 4 holds if we weaken the definition of compatibility of collections of disks to insist on angle compatibility $\measuredangle(D_i, D_j) = \measuredangle(\tilde{D}_i, \tilde{D}_j)$ if and only if one of $D_i \cap D_j$ and $\tilde{D}_i \cap \tilde{D}_j$ contains the other. Still, if we hope to remove the thinness condition, then it seems very likely that angle compatibility will be essential.

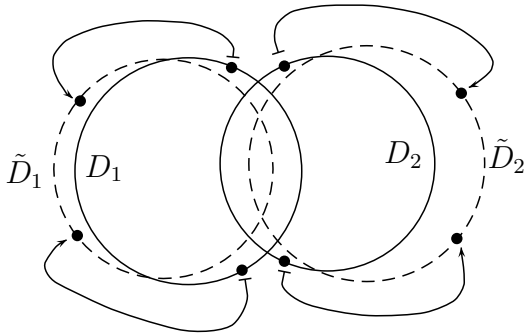

Figure 3.15: **A counterexample to Theorem 4 if we allow $\measuredangle(D_1, D_2) \neq \measuredangle(\tilde{D}_1, \tilde{D}_2)$.** Any indexable $f : \partial(D_1 \cup D_2) \to \partial(\tilde{D}_1 \cup \tilde{D}_2)$ making the shown identifications gives $\eta(f) = -1$.

We state a precise conjecture to end the chapter. Suppose that $\mathcal{K} = \{K_1, \ldots, K_n\}$ and $\tilde{\mathcal{K}} = \{\tilde{K}_1, \ldots, \tilde{K}_n\}$ are collections of convex closed Jordan domains having smooth boundaries, so that there are no $i \neq j$ so that one of K_i and K_j contains the other, or that one of $\tilde{K}_i$ and $\tilde{K}_j$ contains the other. We say that $\mathcal{K}$ and $\tilde{\mathcal{K}}$ are *compatible* if the following conditions hold.

- For every i we have that K_i and $\tilde{K}_i$ are positive homothetic images of one another, that is, there exists a positive real number a and a $b \in \mathbb{C}$ so that $\tilde{\mathcal{K}}$ is the image of K under $z \mapsto az + b$.

- The sets K_i and K_j meet if and only if the sets $\tilde{K}_i$ and $\tilde{K}_j$ meet. Then the sets in one pair have tangent boundaries if and only if the sets in the other pair do.

- Suppose that K_i and K_j overlap, and orient their boundaries as usual. Let $u_{i\to j}$ be the point of $\partial K_i \cap \partial K_j$ where ∂K_i enters K_j. Define $\tilde{u}_{i\to j}$ similarly. Then we insist that the angle between K_i and K_j at $u_{i\to j}$ is the same as the angle between $\tilde{K}_i$ and $\tilde{K}_j$ at $\tilde{u}_{i\to j}$, where the *angle between K_i and K_j* is defined analogously to $\measuredangle(A, B)$ the angle between two disks.

Then we make the following conjecture.

Conjecture 3.14. *Let $\mathcal{K} = \{K_1, \ldots, K_n\}$ and $\tilde{\mathcal{K}} = \{\tilde{K}_1, \ldots, \tilde{K}_n\}$ be compatible collections of convex closed Jordan domains in the plane having smooth boundaries. Let $f : \partial\mathcal{K} \to \partial\tilde{\mathcal{K}}$ be a faithful indexable homeomorphism. Then $\eta(f) \geq 0$.*

It seems that there is no reasonable way to get rid of the convexity condition, as there are examples of non-convex closed Jordan domains K and $\tilde{K}$ with smooth boundaries that are homothetic images of one another, so that ∂K and $\partial\tilde{K}$ meet four times, allowing us to generate counter-examples in the spirit of Figure 3.2.

Chapter 4

Proofs of the main results

In Section 4.1, we prove the Index Theorem 4, modulo four propositions. We give the complete statements of these propositions in our outline, and number them according to where they are found with proofs in the text.

4.1 Proof outline for the Index Theorem 4

For the remainder of this chapter, fix $\mathcal{D} = \{D_1, \ldots, D_n\}$ and $\tilde{\mathcal{D}} = \{\tilde{D}_1, \ldots, \tilde{D}_n\}$, and a faithful indexable homeomorphism $f : \partial\mathcal{D} \to \partial\tilde{\mathcal{D}}$, as in the statement of Theorem 4. Our proof is by induction. The base case is exactly the Circle Index Lemma 3.1, so we suppose from now on that $n \geq 2$.

4.1.1 Extending f to enable induction

We say that two closed disks *overlap* if their interiors meet. Suppose that $D_i \neq D_j$ overlap. Then the *eye* between them is defined to be $E_{ij} = E_{ji} = D_i \cap D_j$. When we quantify over the eyes E_{ij} in $\mathcal{D}$, we keep in mind that $E_{ij} = E_{ji}$ and treat this as a single case. Eyes in $\tilde{\mathcal{D}}$ are defined analogously. A homeomorphism $e_{ij} : \partial E_{ij} \to \partial\tilde{E}_{ij}$ is called *faithful* if it restricts to homeomorphisms $D_i \cap \partial E_{ij} \to \tilde{D}_i \cap \partial\tilde{E}_{ij}$ and $D_j \cap \partial E_{ij} \to \tilde{D}_j \cap \partial\tilde{E}_{ij}$.

Observation 4.1. *For every eye E_{ij} there exists a faithful indexable homeomorphism $e_{ij} : \partial E_{ij} \to \partial\tilde{E}_{ij}$. Furthermore, however they are chosen, the homeomorphisms e_{ij} agree on pairwise intersections of their domains, and every e_{ij} agrees with f on $\partial E_{ij} \cap \partial\mathcal{D}$.*

The only way that there could fail to exist any faithful fixed-point-free homeomorphisms $\partial E_{ij} \to \partial\tilde{E}_{ij}$ is if a pair of corresponding points in $\partial D_i \cap \partial D_j$ and $\partial\tilde{D}_i \cap \partial\tilde{D}_j$ coincide. This cannot happen by the general position hypothesis on $\mathcal{D}$ and $\tilde{\mathcal{D}}$. That the e_{ij} agree with f follows from the faithfulness conditions on f and on the e_{ij}.

The e_{ij} agree with one another by faithfulness and because if three distinct disks D_i, D_j, D_k meet at some point z, then $z \in \partial D_i \cap \partial D_j \cap \partial D_k$.

For every E_{ij} fix a faithful indexable e_{ij}. For $i \in \{1, \ldots, n\}$ let $g_i : \partial D_i \to \partial \tilde{D}_i$ be the function induced by restricting to f or to the e_{ij}, as necessary. It is routine to check that g_i defined this way is an indexable homeomorphism. The following observation serves as a good intuition builder, and will be appealed to in the induction step of our proof.

Observation 4.2. $\eta(f) = \sum_{i=1}^{n} \eta(g_i) - \sum_{E_{ij}} \eta(e_{ij})$

This follows from the Index Additivity Lemma 3.2: notice that $\eta(e_{ij})$ is exactly double-counted in the sum $\eta(g_i) + \eta(g_j)$.

If $I \subset \{1, \ldots, n\}$ then let $\mathcal{D}_I = \{D_i : i \in I\}$, similarly $\tilde{\mathcal{D}}_I$. We denote by $f_I : \partial \mathcal{D}_I \to \partial \tilde{\mathcal{D}}_I$ the function obtained by restriction to f or to the e_{ij}, as necessary. Then f_I is a faithful indexable homeomorphism. We make another observation.

Observation 4.3. *Let $I, J \subset \{1, \ldots, n\}$ be disjoint and non-empty, satisfying $I \sqcup J = \{1, \ldots, n\}$. Then by fixed-point index additivity we get*

$$\eta(f) = \eta(f_I) + \eta(f_J) - \sum \eta(e_{ij})$$

where the sum is taken over all E_{ij} so that $i \in I, j \in J$.

4.1.2 The induction step

We first make a simplifying observation that gives us access to our main propositions:

Observation 4.4. *Suppose that $D_j \setminus \cup_{i \neq j} D_i$ and $\tilde{D}_j \setminus \cup_{i \neq j} \tilde{D}_i$ are disjoint for some j. Then we are done by induction.*

Proof. If neither of D_j and $\tilde{D}_j$ contains the other, then j does not belong to any subsumptive subset of $\{1, \ldots, n\}$, so we are done by applying induction. Thus suppose without loss of generality that $\tilde{D}_j \subset D_j$. Then there must be an $i \neq j$ so that $\tilde{D}_j \subset D_i$. Let J be the maximal subsumptive subset of $\{1, \ldots, n\}$ containing j. If $\tilde{D}_i \not\subset D_i$, then $i \notin J$, but $\tilde{E}_{ij} \subset E_{ij}$, so J is not isolated. Furthermore, one can show that in this case $J = \{j\}$, and we are done by induction. Thus suppose that $\tilde{D}_i \subset D_i$, so $i \in J$. Then because $\tilde{D}_j \subset D_i$ one can show that J is isolated in $\{1, \ldots, n\}$ if and only if $J \setminus \{j\}$ is isolated in $\{1, \ldots, n\} \setminus \{j\}$, and we are done by induction. $\square$

Thus we may assume without loss of generality for the remainder of the chapter that $D_j \setminus D_i$ and $\tilde{D}_j \setminus \tilde{D}_i$ meet, for all i, j.

The following proposition will be key in our induction step:

Proposition 6.7. *Let $\{A, B\}$ and $\{\tilde{A}, \tilde{B}\}$ be pairs of overlapping closed disks in the plane $\mathbb{C}$ in general position. Suppose that neither of $E = A \cap B$ and $\tilde{E} = \tilde{A} \cap \tilde{B}$ contains the other. Suppose further that $A \setminus B$ and $\tilde{A} \setminus \tilde{B}$ meet, and that $B \setminus A$ and $\tilde{B} \setminus \tilde{A}$ meet. Then there is a faithful indexable homeomorphism $e : \partial E \to \partial \tilde{E}$ satisfying $\eta(e) = 0$.*

Thus for example if for no i, j is it the case that one of E_{ij} and $\tilde{E}_{ij}$ contains the other, then we are done by induction. Alternatively, if there exist disjoint non-empty $I, J \subset \{1, \ldots, n\}$ so that $I \sqcup J = \{1, \ldots, n\}$, and so that for every $i \in I, j \in J$ we have that neither of E_{ij} and $\tilde{E}_{ij}$ contains the other, then again we are done by induction. For the remainder of the section, suppose we have fixed faithful indexable e_{ij} so that $\eta(e_{ij}) = 0$ whenever neither of E_{ij} and $\tilde{E}_{ij}$ contains the other, and so that $\eta(e_{ij}) = 1$ otherwise. Next:

Proposition 4.10. *Let $\mathcal{D} = \{D_1, \ldots, D_n\}$ and $\tilde{\mathcal{D}} = \{\tilde{D}_1, \ldots, \tilde{D}_n\}$ be as in the statement of Theorem 4. Let I be a maximal non-empty subsumptive subset of $\{1, \ldots, n\}$. Then there is at most one pair $i \in I$, $j \in \{1, \ldots, n\} \setminus I$ so that one of $E_{ij} = D_i \cap D_j$ and $\tilde{E}_{ij} = \tilde{D}_i \cap \tilde{D}_j$ contains the other.*

Note that *a priori*, even if $I \subset \{1, \ldots, n\}$ is subsumptive, then it may still be that $\eta(f_I) \neq 1$, for example see Figure 4.4 on p. 86. Fortunately, the following holds:

Proposition 4.9. *Let $n \geq 3$ be an integer. Let $\{D_i : i \in \mathbb{Z}/n\mathbb{Z}\}$ and $\{\tilde{D}_i : i \in \mathbb{Z}/n\mathbb{Z}\}$ be collections of closed disks in the plane $\mathbb{C}$ in general position so that the following conditions hold.*

- *We have that $\tilde{D}_i$ is contained in the interior of D_i for all i.*
- *The disk D_i overlaps with $D_{i\pm1}$, and the disk $\tilde{D}_i$ overlaps with $\tilde{D}_{i\pm1}$, for all i.*
- *The disks D_{i-1} and D_{i+1} do not meet in the interior of D_i for any i, similarly for $\tilde{D}_{i-1}, \tilde{D}_{i+1}, \tilde{D}_i$.*

Then $\sum_{i \in \mathbb{Z}/n\mathbb{Z}} \measuredangle(\tilde{D}_i, \tilde{D}_{i+1}) < \sum_{i \in \mathbb{Z}/n\mathbb{Z}} \measuredangle(D_i, D_{i+1})$. In particular it cannot be that $\measuredangle(D_i, D_j) = \measuredangle(\tilde{D}_i, \tilde{D}_j)$ for all i, j.

Thus if $I \subset \{1, \ldots, n\}$ is subsumptive then $\cup_{i \in I} D_i$ and $\cup_{i \in I} \tilde{D}_i$ are compact Jordan domains. It follows by the Circle Index Lemma 3.1 that $\eta(f_I) = 1$. In this case we are done by induction and Observation 4.3. Thus we may suppose for the remainder of the proof, without loss of generality, that there is no i so that one of D_i and $\tilde{D}_i$ contains the other.

We now state our final technical proposition:

Proposition 6.11. *Let $\mathcal{D} = \{D_1, \ldots, D_n\}$ and $\tilde{\mathcal{D}} = \{\tilde{D}_1, \ldots, \tilde{D}_n\}$ be as in the statement of Theorem 4, and so that for all i, j the sets $D_i \setminus D_j$ and $\tilde{D}_i \setminus \tilde{D}_j$ meet. Suppose that there is no i so that one of D_i and $\tilde{D}_i$ contains the other. Suppose that for every pair of disjoint non-empty subsets $I, J \subset \{1, \ldots, n\}$ so that $I \sqcup J = \{1, \ldots, n\}$, there exists an eye E_{ij} with $i \in I$ and $j \in J$ so that one of E_{ij} and $\tilde{E}_{ij}$ contains the other. Then for every i we have that any faithful indexable homeomorphism $g_i : \partial D_i \to \partial \tilde{D}_i$ satisfies $\eta(g_i) \geq 1$. Furthermore there is a k so that D_i and D_k overlap for all i, and so that one of E_{ij} and $\tilde{E}_{ij}$ contains the other if and only if either $i = k$ or $j = k$.*

Unless one of our earlier propositions has finished off the proof for us by induction, the hypotheses of Proposition 6.11 hold, and we are done by Observation 4.2. Thus Theorem 4 will be proved once we establish Propositions 4.9, 4.10, 6.7, and 6.11. □

We establish Propositions 4.9 and 4.10 in Section 4.2. This is because their proofs are quick and elementary, and some ingredients of their proofs are needed for our proof of our main rigidity theorems. We then prove our main rigidity theorems in Section 4.3. Afterward, we outline the rest of the paper in Section 4.4.

4.2 Subsumptive collections

In this section we prove Propositions 4.9 and 4.10 having to do with subsumptive subsets of the index set of our collections of disks. Before doing so, we first establish some key geometric facts in Section 4.2.1.

4.2.1 Key geometric lemmas

We first make the following observation:

Observation 4.5. *Suppose that A_1, A_2, A_3, A_4 are four closed disks so that $\mathbb{C} \setminus \cup_{i=1}^4 A_i$ has a bounded component U whose boundary is a curvilinear quadrilateral, with every A_i contributing a side. Suppose the A_i are labeled so that as we traverse ∂U we arrive at the A_i in cyclic order. Then $\sum_{i=1}^4 \measuredangle(A_i, A_{i+1})$, denoting $A_5 = A_1$, is strictly less than 2π. See Figure 4.1.*

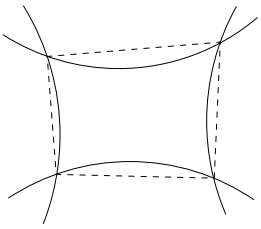

Figure 4.1: **A complementary component of the union of four closed disks having a curvilinear quadrilateral as its boundary.** The sum of the angles inside of the dashed honest quadrilateral is exactly 2π. This sum is greater than the sum of the shown external intersection angles of the disks.

We use Observation 4.5 to prove the following key lemma:

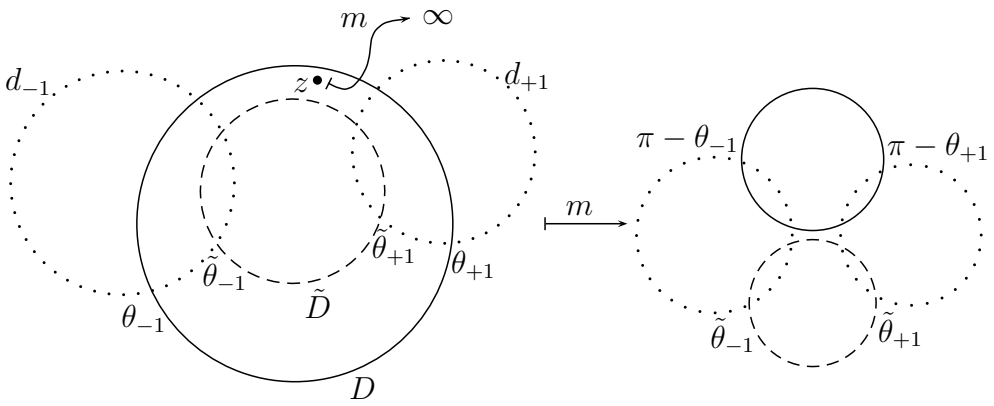

Figure 4.2: **A Möbius transformation chosen to prove Lemma 4.6.** On the right side of the figure note that the angles inside of the small curvilinear quadrilateral in the middle of the four disks are exactly the external intersection angles of the disks.

Lemma 4.6. *Let $d_{-1}, d_{+1}, D, \tilde{D}$ be closed disks in $\mathbb{C}$, so that $\tilde{D}$ is contained in the interior of D, so that both of D and $\tilde{D}$ meet both of d_{-1} and d_{+1}, and so that the interiors of d_{-1}, d_{+1}, and D do not meet. Suppose that neither of d_{-1} and d_{+1} is contained in D. We will denote $\theta_{-1} = \measuredangle(D, d_{-1})$, similarly θ_{+1}, $\tilde{\theta}_{-1}$, $\tilde{\theta}_{+1}$. Then $\tilde{\theta}_{-1} + \tilde{\theta}_{+1} < \theta_{-1} + \theta_{+1}$.*

Proof. Suppose that d_{-1} and d_{+1} are disjoint, as in Figure 4.2. Let z be a point in the interior of $D \setminus (\tilde{D} \cup d_{-1} \cup d_{+1})$, and let m be a Möbius transformation sending z to ∞. Then m inverts the disk D but none of the disks $\tilde{D}, d_{-1}, d_{+1}$. Because m preserves angles we get $(\pi - \theta_{-1}) + (\pi - \theta_{+1}) + \tilde{\theta}_{-1} + \tilde{\theta}_{+1} < 2\pi$ by Observation 4.5, and the desired inequality follows. We leave the case where d_{-1} and d_{+1} meet to the reader. The same proof works. The key ingredient is that $D \setminus (\tilde{D} \cup d_{-1} \cup d_{+1})$ has several connected components, one of which is a curvilinear quadrilateral with the

appropriate interior angles. Then we choose the point z to lie in one of the other components. □

The following is easy to prove as a corollary to Lemma 4.6, by applying a suitable Möbius transformation:

Lemma 4.7. *Let A, B be a pair of closed disks in the plane $\mathbb{C}$ which overlap, similarly $\tilde{A}, \tilde{B}$, so that $\measuredangle(A, B) = \measuredangle(\tilde{A}, \tilde{B})$. Suppose that $\tilde{A}$ is contained in the interior of A and that $\tilde{B}$ is contained in the interior of B. Suppose also that neither $\tilde{A} \subset B$ nor $\tilde{B} \subset A$. Then $2\measuredangle(A, B) = 2\measuredangle(\tilde{A}, \tilde{B}) < \measuredangle(\tilde{A}, B) + \measuredangle(A, \tilde{B})$.*

For the next lemma, see Figure 4.3:

Lemma 4.8. *Let A, B, C be closed disks, none of which is contained in any other. Suppose that A and C overlap, and $C \cap A \subset B$. Then $\measuredangle(A, C) < \measuredangle(A, B)$.*

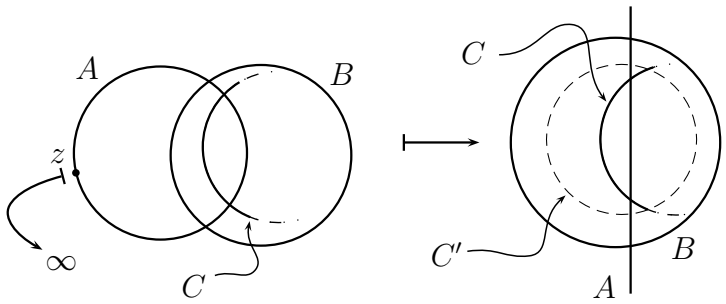

Figure 4.3: **A Möbius transformation chosen to prove Lemma 4.8.**

Proof. Let $z \in \partial A \setminus B$. Because of the hypothesis that $A \cap C \subset B$, we have that $z \notin C$. Apply a Möbius transformation sending $z \mapsto \infty$ so that A becomes the left half-plane. Because $z \notin B, C$ we have that B and C remain closed disks after this transformation. Let C' be the closed disk so that $\measuredangle(A, C') = \measuredangle(A, B)$, and so that C and C' have the same Euclidean radius and the same vertical Euclidean coordinate. Then C is obtained from C' by a translation to the right or to the left. In fact it must be a translation to the right, because the points $\partial B \cap \partial C$ must lie in the complement of A, which is the right-half plane. But it is easy to see that $\measuredangle(A, C')$ is monotone decreasing as C' slides to the right. □

4.2.2 There cannot be a subsumptive loop of disks

For our next proposition, see Figure 4.4 for an example.

Proposition 4.9. *Let $n \geq 3$ be an integer. Let $\{D_i : i \in \mathbb{Z}/n\mathbb{Z}\}$ and $\{\tilde{D}_i : i \in \mathbb{Z}/n\mathbb{Z}\}$ be collections of closed disks in the plane $\mathbb{C}$ in general position so that the following conditions hold.*

- *We have that $\tilde{D}_i$ is contained in the interior of D_i for all i.*
- *The disk D_i overlaps with $D_{i\pm1}$, and the disk $\tilde{D}_i$ overlaps with $\tilde{D}_{i\pm1}$, for all i.*
- *The disks D_{i-1} and D_{i+1} do not meet in the interior of D_i for any i, similarly for $\tilde{D}_{i-1}, \tilde{D}_{i+1}, \tilde{D}_i$.*

Then $\sum_{i\in\mathbb{Z}/n\mathbb{Z}} \measuredangle(\tilde{D}_i, \tilde{D}_{i+1}) < \sum_{i\in\mathbb{Z}/n\mathbb{Z}} \measuredangle(D_i, D_{i+1})$. In particular it cannot be that $\measuredangle(D_i, D_j) = \measuredangle(\tilde{D}_i, \tilde{D}_j)$ for all i, j.

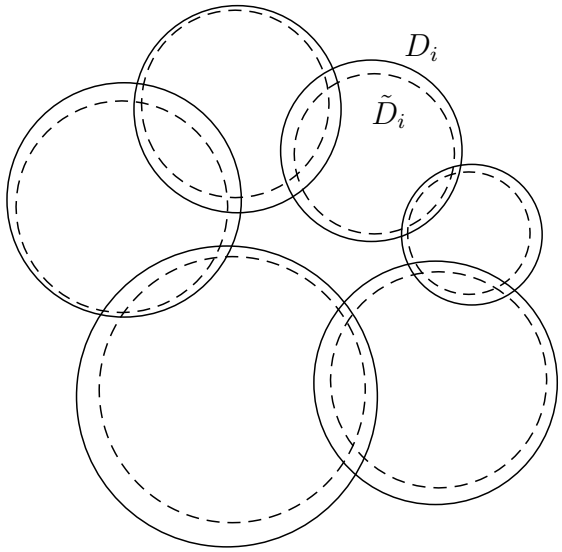

Figure 4.4: **Two closed chains of disks with $\tilde{D}_i \subsetneq D_i$ for all i.** The solid disks are the D_i and the dashed disks are the $\tilde{D}_i$. Proposition 4.9 implies that $\measuredangle(\tilde{D}_i, \tilde{D}_{i+1}) \neq \measuredangle(D_i, D_{i+1})$ for some i.

Proof. Note first that for $\measuredangle(D_i, D_{i+1})$ to be well-defined, we need to show that neither $D_i \subset D_{i+1}$ nor $D_{i+1} \subset D_i$. The same is true for $\measuredangle(\tilde{D}_i, \tilde{D}_{i+1})$. Suppose for contradiction that $D_i \subset D_{i+1}$. Then D_{i-1} and D_i overlap inside of D_{i+1}, contradicting our hypotheses. By symmetry we get that $D_{i+1} \not\subset D_i$, and the two disks overlap by hypothesis. The proof that $\measuredangle(\tilde{D}_i, \tilde{D}_{i+1})$ is well-defined is identical.

To finish off the proof, we apply Lemma 4.6 twice. In both cases we will let $D = D_i$ and $\tilde{D} = \tilde{D}_i$. First let $d_{-1} = D_{i-1}$ and $d_{+1} = D_{i+1}$. This gives:

$$\measuredangle(D_{i-1}, \tilde{D}_i) + \measuredangle(D_{i+1}, \tilde{D}_i) < \measuredangle(D_{i-1}, D_i) + \measuredangle(D_{i+1}, D_i) \tag{4.1}$$

Next let $d_{-1} = \tilde{D}_{i-1}$ and $d_{+1} = \tilde{D}_{i+1}$. This gives:

$$\measuredangle(\tilde{D}_{i-1}, \tilde{D}_i) + \measuredangle(\tilde{D}_{i+1}, \tilde{D}_i) < \measuredangle(D_i, \tilde{D}_{i-1}) + \measuredangle(D_i, \tilde{D}_{i+1}) \tag{4.2}$$

Now note that if we let i range over $\mathbb{Z}/n\mathbb{Z}$, the sum of the terms on the left side of equation 4.1 is equal to the sum of the terms on the right side of equation 4.2. The desired inequality follows. □

4.2.3 Mostly-isolation of subsumptive collections

Proposition 4.10. *Let $\mathcal{D} = \{D_1, \ldots, D_n\}$ and $\tilde{\mathcal{D}} = \{\tilde{D}_1, \ldots, \tilde{D}_n\}$ be as in the statement of Theorem 4, thin and in general position. Suppose there is some D_i so that $\tilde{D}_i \subset D_i$. Let I be a maximal subsumptive subset of $\{1, \ldots, n\}$. Then there is at most one pair $i \in I$, $j \in \{1, \ldots, n\} \setminus I$ so that D_i and D_j overlap and one of $E_{ij} = D_i \cap D_j$ and $\tilde{E}_{ij} = \tilde{D}_i \cap \tilde{D}_j$ contains the other.*

Recall that $I \subset \{1, \ldots, n\}$ is *subsumptive* if $\cup_{i \in I} D_i$ is connected and either $\tilde{D}_i \subset D_i$ for all $i \in I$ or $D_i \subset \tilde{D}_i$ for all $i \in I$. The rest of the section is spent proving Proposition 4.10. Fix I as in the statement of Proposition 4.10, and suppose without loss of generality that $\tilde{D}_i \subset D_i$ for all $i \in I$.

First, let G_{u} be the undirected simple graph defined as follows: the vertex set is I, and there is an edge between i and j if and only if D_i and D_j overlap. Observe:

Observation 4.11. *The graph G_{u} is connected and is a tree.*

This follows from Proposition 4.9 and the general position hypothesis.

Next, let G be the directed graph so that $\langle i \to j \rangle$ is an edge of G if and only if:

- we have that $\langle i, j \rangle$ is an edge of G_{u}, and
- either $\measuredangle(\tilde{D}_i, D_j) > \measuredangle(D_i, D_j)$ or $\tilde{D}_i \subset D_j$.

If $\langle i \to j \rangle$ is an edge of G then we call $\langle i \to j \rangle$ an edge *pointing away from i in G*. The idea is that if $\langle i \to j \rangle$ is an edge in G then the disk $\tilde{D}_i \subset D_i$ is "shifted towards D_j in D_i." See Figure 4.5 for an example. We make a series of observations about G. We advise the reader to "verify" each by an examination of Figure 4.5, or by trying to draw a counter-example, for the purpose of building intuition.

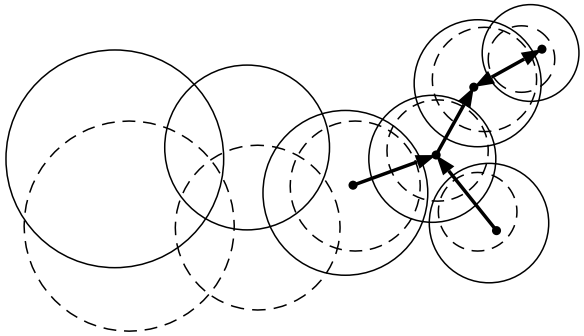

Figure 4.5: **The directed graph G associated to a maximal subsumptive I.** The solid disks are the D_i and the dashed disks are the $\tilde{D}_i$. The graph G_{u} can be obtained by undirecting every edge.

Observation 4.12. *If $\langle i, j\rangle$ is an edge in G_{u} then at least one of $\langle i \to j\rangle$ and $\langle j \to i\rangle$ is an edge in G, and possibly both are.*

This follows from Lemma 4.7.

Observation 4.13. *If $i \in I$ then there is at most one edge $\langle i \to j\rangle$ in G pointing away from i.*

This follows from Lemma 4.6, with $D = D_i$, $\tilde{D} = \tilde{D}_i$, $d_{-1} = D_j$, $d_{+1} = D_k$, for $j, k \in I$ so that D_i overlaps with both D_j and D_k.

Observation 4.14. *Let $\langle i_1, i_2, \ldots, i_m\rangle$ be a simple path in G_{u}, meaning that $\langle i_\ell, i_{\ell+1}\rangle$ is an edge in G_{u} for all $1 \leq \ell < m$ and that i_ℓ and $i_{\ell'}$ are distinct for $\ell \neq \ell'$. Suppose that $\langle i_{m-1} \to i_m\rangle$ is an edge in G. Then $\langle i_\ell \to i_{\ell+1}\rangle$ is an edge in G for $1 \leq \ell < m$.*

This follows from Observations 4.12 and 4.13, and induction.

Observation 4.15. *There is at most one $i \in I$ so that there is no edge pointing away from i in G.*

This follows from Observations 4.12, 4.13, and 4.14. If there is an i as in the statement of Observation 4.15, then we call this i the *sink* of the directed tree G.

Having established all we need to about G, we are ready to make two final observations which will complete the proof of Proposition 4.10. First:

Observation 4.16. *Let $i \in I$. Then there is at most one $1 \leq j \leq n$ different from i so that D_i and D_j overlap and either $\tilde{D}_i \subset D_j$ or $\measuredangle(D_i, D_j) < \measuredangle(\tilde{D}_i, D_j)$.*

This follows from Lemma 4.6 in the same way as does Observation 4.13. Next:

Observation 4.17. *Suppose that i and j are as in the statement of Proposition 4.10. That is, we have that $i \in I$ and $j \neq I$ so that D_i and D_j overlap and $\tilde{E}_{ij} \subset E_{ij}$. Suppose $\tilde{D}_i \subset D_i$. Then $\measuredangle(D_i, D_j) = \measuredangle(\tilde{D}_i, \tilde{D}_j) < \measuredangle(\tilde{D}_i, D_j)$.*

This follows by an application of Lemma 4.8 with $\tilde{D}_i = A$, $D_j = B$, and $\tilde{D}_j = C$. Thus if i and j are as in the statement of Proposition 4.10, then i is the unique sink of G. Furthermore by Observations 4.16 and 4.17 there is no $k \in \{1, \ldots, n\} \setminus I$ different from j so that D_i and D_k overlap and so that one of E_{ik} and $\tilde{E}_{ik}$ contains the other. Proposition 4.10 follows. □

4.3 Proof of our main rigidity theorems

The main idea of the proofs here is very similar to that of the proof of Theorem 3.7 given in Section 3.5, but with extra issues to deal with. Therefore we recommend for the reader to review that section now. We will prove only Rigidity Theorem 1. Adapting the argument to prove Rigidity Theorem 2 and Uniformization Theorem 3 is straightforward, proceeding identically to the way that we adapted our proof of Theorem 3.7 to prove Theorems 3.8 and 3.9.

Let $\mathcal{C}$ and $\tilde{\mathcal{C}}$ be as in the statement of Rigidity Theorem 1. That is, both are thin disk configurations realizing (G, Θ), where $G = (V, E)$ is a mostly-triangulation of the 2-sphere $\mathbb{S}^2$ and $\Theta : E \to [0, \pi)$. We wish to show that $\mathcal{C}$ and $\tilde{\mathcal{C}}$ differ by a Möbius transformation. As in the proof of Theorem 3.8, our goal here is to superimpose the configurations $\mathcal{C}$ and $\tilde{\mathcal{C}}$ on the Riemann sphere in a convenient way. In particular, we wish to find two superimposed topological quadrilaterals Q and $\tilde{Q}$ so that any indexable homeomorphism $\partial Q \to \partial \tilde{Q}$ identifying corresponding corners has a negative fixed-point index, as in Figure 3.8b. There are more cases to handle than in the proof of Theorem 3.8, and the discussion of how to superimpose $\mathcal{C}$ and $\tilde{\mathcal{C}}$ makes up the first half of the proof, contained in Section 4.3.1. After this is done, we as before construct an indexable homeomorphism $\partial Q \to \partial \tilde{Q}$ identifying corresponding corners having a non-negative fixed-point index, obtaining a contradiction. This is where Theorem 4 comes in. The construction is more involved here than in the proof of Theorem 3.8, and makes up the second half of the proof, contained in Section 4.3.2.

4.3.1 Finding the quadrilaterals Q and $\tilde{Q}$

We work on the Riemann sphere $\hat{\mathbb{C}}$. Let $\psi : G \to \mathbb{S}^2$ be a mostly-embedding of G, and let F be the faces of the cell decomposition of $\mathbb{S}^2$ obtained by deleting every pair of edges e_1, e_2 so that $\psi(e_1)$ and $\psi(e_2)$ cross. Then F consists of triangles and

quadrilaterals.

Pick a face $f_0 \in F$. First suppose that $f_0 = \langle v_1, v_2, v_3 \rangle$ is a triangular face. If the disks D_1, D_2, D_3, equivalently $\tilde{D}_1, \tilde{D}_2, \tilde{D}_3$, form an interstice, then normalize both configurations $\mathcal{C}$ and $\tilde{\mathcal{C}}$ so that $\infty \in \hat{\mathbb{C}}$ lies in the interstices. Otherwise, for both triples the three constituent disks meet at a single point. In that case, normalize both configurations so that this point is ∞. Next, suppose that $f_0 = \langle v_0, v_1, v_2, v_3 \rangle$ is a quadrilateral face. Then the disks D_1, D_2, D_3, D_4 meet at a single point. In this case, normalize both configurations so that this point is ∞. We will apply further normalizations later. For now, the possibilities for $\mathcal{C}$ are shown in Figure 4.6.

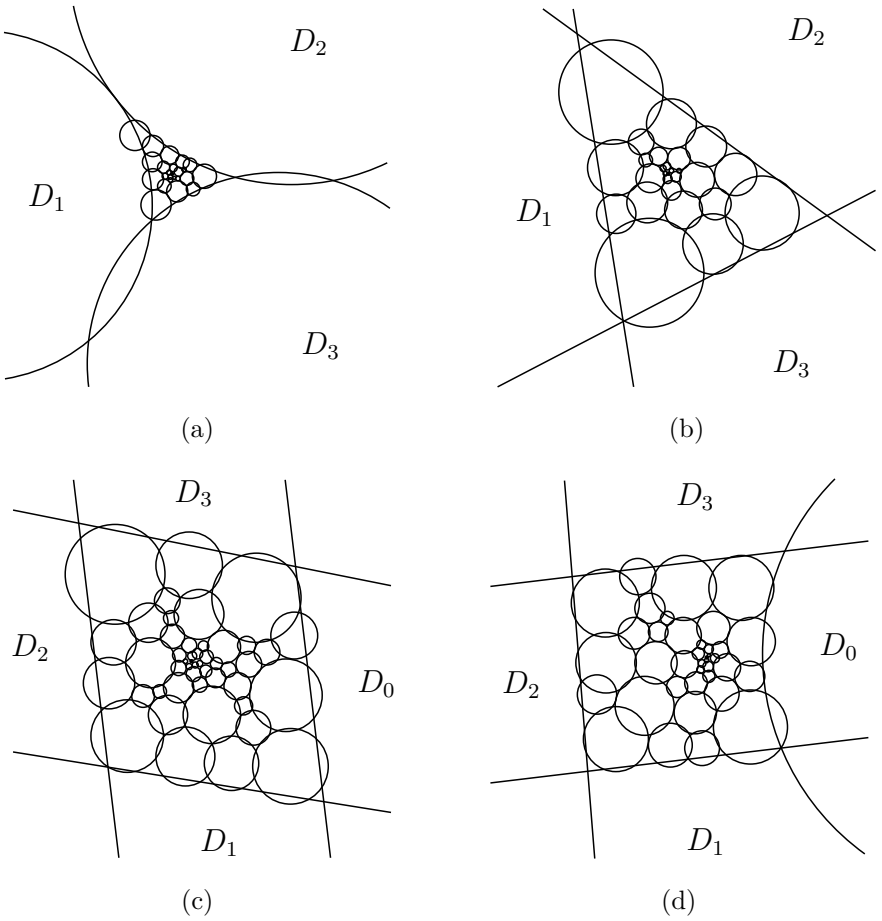

Figure 4.6: **The possibilities for $\mathcal{C}$ after an initial normalization.** We get (a) if and only if D_1, D_2, D_3 form an interstice. Otherwise we get one of (b), (c), (d).

Next, regardless of whether f_0 is triangular or quadrilateral, normalize so that D_i and $\tilde{D}_i$ coincide for $i = 1, 2, 3$. We are supposing for contradiction that $\mathcal{C}$ and $\tilde{\mathcal{C}}$ are not Möbius equivalent. The following is along the lines of Observation 3.10:

Observation 4.18. *We may suppose without loss of generality that there is a pair of*

disks D_4 and $\tilde{D}_4$ having different spherical radii on $\hat{\mathbb{C}}$, so that either v_4 belongs to a face sharing an edge with f_0, or f_0 is a quadrilateral and $v_4 = v_0$.

We now enumerate the possibilities, and in each case describe how we obtain Q and $\tilde{Q}$ in terms of one positive real parameter ε. From now on we work in the plane, in the sense that all further normalizations and transformations are Euclidean similarities of the plane. There are four cases:

Case 1. The face $f_0 = \langle v_1, v_2, v_3 \rangle$ is triangular, and the disks D_1, D_2, D_3, equivalently the disks $\tilde{D}_1, \tilde{D}_2, \tilde{D}_3$, form an interstice.

Here Figure 4.6a occurs. This case is very similar to the situation in the proof of Theorem 3.8. We may suppose without loss of generality that v_4 belongs to a face f_1, which may be quadrilateral or triangular, sharing the edge $\langle v_2, v_3 \rangle$ with f_0. Then let $V_Q = V \setminus \{v_1, v_2, v_3, v_4\}$ and $F_Q = F \setminus \{f_0, f_1\}$. As before let U_f denote the interstice corresponding to the face $f \in F$, if it exists. Then define:

$$Q = \bigcup_{v \in V_Q} D_v \cup \bigcup_{f \in F_Q} U_f \setminus \bigcup E_{ij}$$

where the last union is taken over E_{ij} so that $v_i \in V_Q$ and $v_j \in V \setminus V_Q$. Define $\tilde{Q}$ analogously.

Rotate and translate so that both intersection points $\partial D_2 \cap \partial D_3$ lie on the horizontal axis, or so that D_2 and D_3 are both tangent to the horizontal axis if they are tangent to each other. Rotate again if necessary to ensure that D_1 lies to the left of $\partial D_2 \cap \partial D_3$. Apply the same motions to $\tilde{\mathcal{C}}$ so that D_i and $\tilde{D}_i$ continue to coincide for $i = 1, 2, 3$. We may suppose without loss of generality that D_4 is bigger than $\tilde{D}_4$. Then a translation to the right by a small amount ε will position our Q and $\tilde{Q}$ where we want them, very similarly to the way shown in Figure 3.8. Let T_ε denote this translation.

It is not as easy to see, but this works even if D_4 and $\tilde{D}_4$ are both tangent to one of $D_2 = \tilde{D}_2$ and $D_3 = \tilde{D}_3$. This case is shown in Figure 4.7. The key point is that after applying T_ε to $\mathcal{C}$, the disk $\tilde{D}_4$ is contained in the interior of D_4.

Case 2. The face $f_0 = \langle v_1, v_2, v_3 \rangle$ is triangular, and the disks D_1, D_2, D_3, equivalently the disks $\tilde{D}_1, \tilde{D}_2, \tilde{D}_3$, meet at a single point which is not a tangency point of any two of them.

Here Figure 4.6b occurs. This case is essentially the same as case 1, except that three of the boundary segments of Q are straight rather than circular, similarly for $\tilde{Q}$. Thus we do not show the figures for this case.

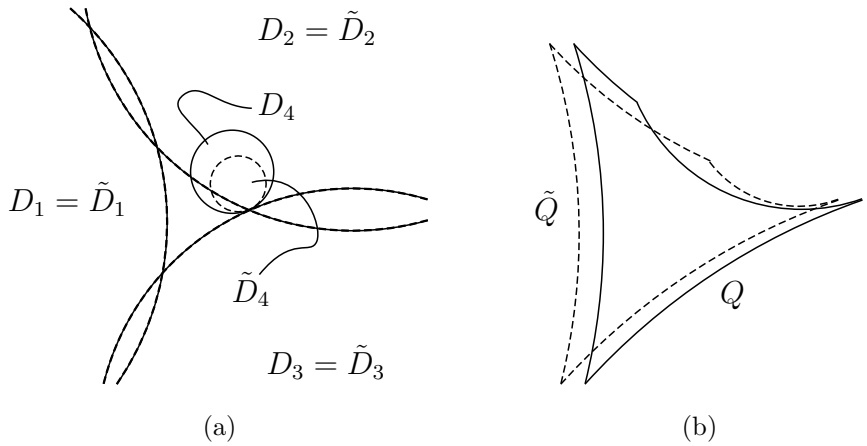

Figure 4.7: **The interactions between $\mathcal{C}$ and $\tilde{\mathcal{C}}$ before and after applying T_ε to $\mathcal{C}$, when Figure 4.6a occurs.**

Case 3. The face $f_0 = \langle v_0, v_1, v_2, v_3 \rangle$ is a quadrilateral.

Here Figure 4.6d occurs. Note that after our normalizations, because D_0 is tangent to D_2, the spherical radius of D_0 corresponds to the distance in the plane between the half-planes D_2 and D_0. The same holds for the disks $\tilde{D}_0$ and $\tilde{D}_2$. Thus the half-planes D_0 and $\tilde{D}_0$ coincide if and only if the spherical radii of the disks D_0 and $\tilde{D}_0$ coincide. We get two natural sub-cases:

Sub-case 3.1. The disks D_0 and $\tilde{D}_0$ have different spherical radii.

Thus we may as well set $v_4 = v_0$. Let $V_Q = V \setminus \{v_1, v_2, v_3, v_4\}$ and $F_Q = F \setminus \{f_0\}$. Then define Q in terms of F_Q and V_Q as before, and define $\tilde{Q}$ analogously.

Then Q and $\tilde{Q}$ are parallelograms. In this case we do not apply a translation to $\mathcal{C}$, but rather a dilation. As before we may suppose without loss of generality that $D_4 = D_0$ has a larger spherical radius than does $\tilde{D}_4 = \tilde{D}_0$, which implies that the half-plane $\tilde{D}_4$ is contained in the half-plane D_4. Then transform $\mathcal{C}$ by a dilation by a factor of $1 + \varepsilon$ about a point contained in the interior of $D_2 \setminus \bigcup_{i \neq 2}(D_i \cup \tilde{D}_i)$. Let T_ε denote this transformation. The situation is shown in Figure 4.8.

Sub-case 3.2. The disks D_0 and $\tilde{D}_0$ have the same spherical radii.

In this case we may suppose without loss of generality that v_4 belongs to a face f_1, which may be quadrilateral or triangular, which shares the edge $\langle v_2, v_3 \rangle$ with f. Then set $V_Q = V \setminus \{v_0, v_1, v_2, v_3, v_4\}$, $F_Q = F \setminus \{f_0, f_1\}$, and define Q and $\tilde{Q}$ as usual.

Normalize $\mathcal{C}$ by a rotation and translation sending the intersection point of ∂D_2 and ∂D_3 to the origin, and sending an arbitrarily chosen point in the interior of

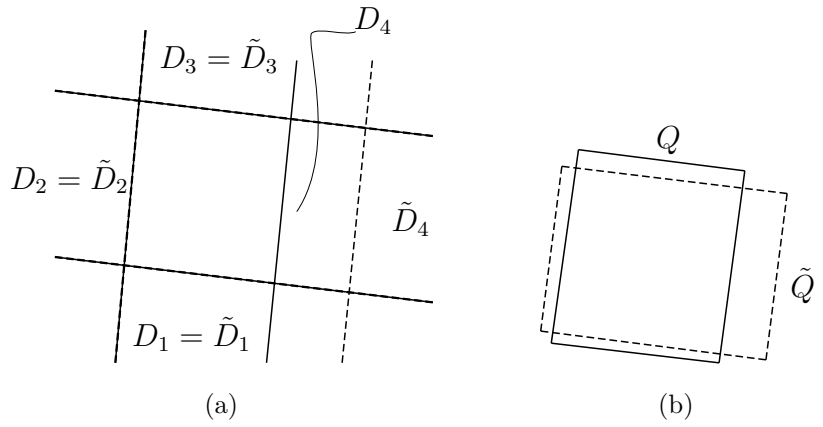

Figure 4.8: **The interactions between $\mathcal{C}$ and $\tilde{\mathcal{C}}$ before and after applying T_ε to $\mathcal{C}$, when the first sub-case of Figure 4.6c occurs.**

$D_2 \cap D_3$ to the positive real axis. Apply the same transformations to $\tilde{\mathcal{C}}$ to preserve $D_i = \tilde{D}_i$ for $i = 0, 1, 2, 3$. We may suppose without loss of generality that D_4 is larger than $\tilde{D}_4$. Then a translation to the right by ε will position Q and $\tilde{Q}$ as we wish. Let T_ε denote this translation. The situation is depicted in Figure 4.9.

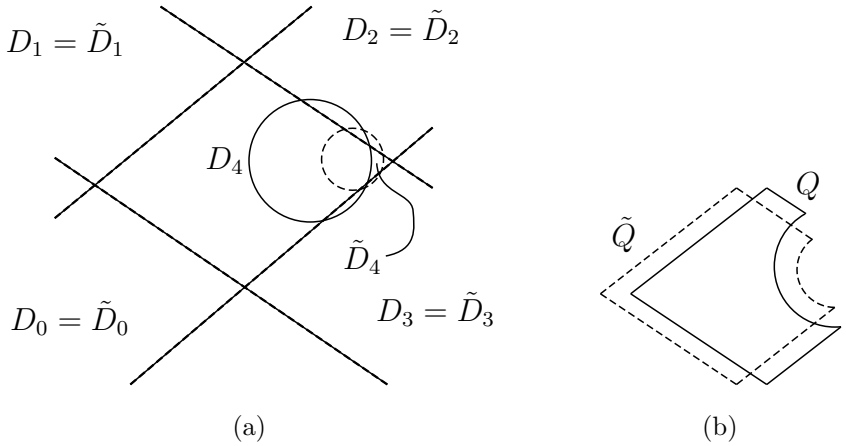

Figure 4.9: **The interactions between $\mathcal{C}$ and $\tilde{\mathcal{C}}$ before and after applying T_ε to $\mathcal{C}$, when the second sub-case of Figure 4.6c occurs.**

Case 4. The face $f_0 = \langle v_1, v_2, v_3 \rangle$ is triangular, and the disks D_1, D_2, D_3, equivalently the disks $\tilde{D}_1, \tilde{D}_2, \tilde{D}_3$, meet at a single point which is a tangency point of two of them.

Here Figure 4.6c occurs. We may suppose without loss of generality that D_2 and D_3 are tangent. Note that because D_2 and D_3 are tangent, the edge $\langle v_2, v_3 \rangle$ cannot belong to a quadrilateral. Thus let v_0 be the vertex other than v_1 which makes a face $\langle v_2, v_3, v_0 \rangle$ with the edge $\langle v_2, v_3 \rangle$. Once we have fixed the Euclidean half-planes D_1, D_2, D_3 via our normalizations, the Euclidean radius of D_0 is determined by the

angle D_0 makes with D_2 and D_3. Thus D_0 and $\tilde{D}_0$ have the same Euclidean radii, because $D_2 = \tilde{D}_2$ and $D_3 = \tilde{D}_3$, and $\tilde{D}_0$ makes the same angles with $\tilde{D}_2$ and $\tilde{D}_3$ as does D_0 with D_2 and D_3 respectively. Instead, the radius of D_0 on $\hat{\mathbb{C}}$ corresponds to the distance of the disk D_0 from the half-plane D_1. We conclude that D_0 and $\tilde{D}_0$ will coincide in the plane if and only if they have the same spherical radii. So, this case is essentially the same as Case 3, and we omit the figures for it.

4.3.2 Obtaining the contradiction via Theorem 4

No matter which case occurred, as before we have that any homeomorphism $\varphi_Q : \partial Q \to \partial \tilde{Q}$ identifying corresponding corners has $\eta(\varphi_Q) = -1$. We wish to construct such a homeomorphism with $\eta(\varphi_Q) \geq 0$ to obtain a contradiction.

Choose ε so that the configurations $\mathcal{C}$ and $\tilde{\mathcal{C}}$ are in general position. For every $f \in F \setminus \{f_0\}$ define $\varphi_f : \partial U_f \to \partial \tilde{U}_f$ as before via the Three Point Prescription Lemma 3.5, to have $\eta(\varphi_f) \geq 0$ and to identify corresponding corners. There may be $f \in F$ so that the corresponding interstices U_f and $\tilde{U}_f$ are empty. We discard and ignore such f.

Suppose that $v_i \in V_Q$ and $v_j \in V \setminus V_Q$ so that D_i and D_j overlap. Denote as usual $E_{ij} = D_i \cap D_j$, similarly $\tilde{E}_{ij}$. Let $\eta(e_{ij})$ be a faithful indexable homeomorphism so that $\eta(e_{ij}) = 0$, unless one of E_{ij} and $\tilde{E}_{ij}$ contains the other, in which case $\eta(e_{ij}) = 1$. We may do this by a slight generalization of Proposition 6.7. Let $\mathcal{D}$ be the set of disks D_v of $\mathcal{C}$ whose vertices v lie in V_Q, similarly $\tilde{\mathcal{D}}$. Then the φ_f and e_{ij} induce a faithful indexable $\varphi_{\mathcal{D}} : \partial \mathcal{D} \to \partial \tilde{\mathcal{D}}$. We get:

$$\eta(\varphi_Q) = \eta(\varphi_{\mathcal{D}}) + \sum_{f \in F_Q} \eta(\varphi_f) - \sum \eta(e_{ij})$$

where the last sum is taken over E_{ij} with $v_i \in V_Q$ and $v_j \in V \setminus V_Q$. We need only to show the following:

Claim 4.19. *If we chose ε small enough, then for every E_{ij} and $\tilde{E}_{ij}$, with $v_i \in V_Q$ and $v_j \in V \setminus V_Q$, so that one of E_{ij} and $\tilde{E}_{ij}$ contains the other, we have that one of D_i and $\tilde{D}_i$ contains the other. Furthermore, the maximal subsumptive subset I of V_Q which contains v_i is isolated, and these subsumptive subsets are distinct.*

This will establish $\eta(\varphi_{\mathcal{D}}) - \sum \eta(e_{ij}) \geq 0$ by Theorem 4, giving $\eta(\varphi_Q) \geq 0$, our desired contradiction.

Proof of Claim 4.19. We break the proof into two cases, the first one relatively easy and the second one slightly more challenging. Suppose without loss of generality that

$\tilde{E}_{ij} \subset E_{ij}$. In both cases we will first show that under a sufficiently small choice of ε we must have that $\tilde{D}_i \subset D_i$, establishing the existence of a subsumptive subset $I \subset V_{\mathcal{D}}$. We then show that, again under a sufficiently small choice of ε, for every $v_k \in V_{\mathcal{D}}$ so that D_k meets $\tilde{D}_i$, we have that $\measuredangle(\tilde{D}_i, D_k) < \measuredangle(D_i, D_k)$. This establishes that I is isolated in $V_{\mathcal{D}}$, and furthermore that v_i is its sink, so these I are distinct.

Case 1. $v_j \neq v_4$

We first consider the picture before applying T_ε to $\mathcal{C}$. Then we have that $D_j = \tilde{D}_j$, so $\measuredangle(\tilde{D}_i, D_j) = \measuredangle(\tilde{D}_i, \tilde{D}_j) = \measuredangle(D_i, D_j)$. In this case it is easy to see that $\tilde{E}_{ij} \subset E_{ij}$ if and only if $\tilde{D}_i \subset D_i$, and that we may ensure that this continues to hold after applying T_ε by picking ε sufficiently small. Furthermore before applying T_ε, if $v_k \in V_Q$ so that D_k meets $\tilde{D}_i$, then we have that $\measuredangle(\tilde{D}_i, D_k) < \measuredangle(D_i, D_k)$, by a slight generalization of Lemma 4.6. Furthermore if ε is sufficiently small then this inequality will persist after we apply T_ε to $\mathcal{C}$.

Case 2. $v_j = v_4$

As in the first case, we consider the picture before applying T_ε to $\mathcal{C}$, and establish our inequalities in this setting, noting that they will persist after applying T_ε if we choose ε sufficiently small. As an example of what can occur, we discuss the case appearing in Figure 4.9. The other cases are handled in essentially the same way. Note first that $\tilde{D}_j = \tilde{D}_4$ is not contained in $D_j = D_4$. It is also clear that if $\tilde{E}_{ij} \subset E_{ij}$ then D_i must meet $\tilde{D}_j$, and furthermore we must have $D_i \cap \tilde{D}_j \subset D_j$. On the other hand, for no pair of disks of $D_i, D_j, \tilde{D}_j$ can it be the case that one contains the other. Then by Lemma 4.8 we get that $\measuredangle(D_i, \tilde{D}_j)$ is strictly smaller than both $\measuredangle(D_i, D_j)$ and $\measuredangle(D_j, \tilde{D}_j)$. It is also clear that if $\tilde{E}_{ij} \subset E_{ij}$ then in particular $\tilde{E}_{ij} \subset D_i$. Then $\tilde{D}_i \subset D_i$, because otherwise we would contradict $\measuredangle(D_i, \tilde{D}_j) < \measuredangle(D_i, D_j) = \measuredangle(\tilde{D}_i, \tilde{D}_j)$ after applying Lemma 4.8 to $A = \tilde{D}_i, B = \tilde{D}_j, C = D_i$.

Finally, if $\tilde{D}_i \subset D_j$ then I is isolated and we are done. Otherwise we will be done by an application of Lemma 4.6 once we show that $\measuredangle(\tilde{D}_i, D_j) > \measuredangle(D_i, D_j) = \measuredangle(\tilde{D}_i, \tilde{D}_j)$. But this follows immediately from an application of Lemma 4.8 to $A = \tilde{D}_i, B = D_j, C = \tilde{D}_i$, assuming $\tilde{D}_i \not\subset D_j$, because $\tilde{D}_i \cap \tilde{D}_j \subset D_j$. This completes the proof of Claim 4.19, and thus of Rigidity Theorem 1. □

4.4 Structure of the rest of the thesis

Two chapters remain, Chapters 5 and 6. In Chapter 5 we establish many topological propositions and lemmas that we will use to fill in the details of the proof

of Theorem 4. First, we establish that when considering the fixed-point index of a homeomorphism $f : \partial K \to \partial \tilde{K}$, we need only to consider the sets K and $\tilde{K}$ up to "simultaneous homeomorphism." Section 5.4 enumerates the relevant "simultaneous homeomorphism" classes of two pairs of overlapping disks $\{A, B\}$ and $\{\tilde{A}, \tilde{B}\}$, essentially reducing the proofs of all remaining propositions to finitely many cases. Section 5.5 proves some technical and unintuitive lemmas that handle special cases not taken care of by our general techniques.

In Chapter 6 we develop a tool, called torus parametrization, that allows us to quickly and easily write down many homeomorphisms $\partial K \to \partial \tilde{K}$ having a desired fixed-point index. We then apply torus parametrization to prove Propositions 6.7 and 6.11, completing the proof of Theorem 4 and thus Theorem 1.

Chapter 5

Topological configurations

Suppose that $X_1, \ldots, X_n$ and $X_1', \ldots, X_n'$ are all subsets of $\mathbb{C}$. Then we say that the collections $\{X_1, \ldots, X_n\}$ and $\{X_1', \ldots, X_n'\}$ are in the same *topological configuration* if there is an orientation-preserving homeomorphism $\varphi : \mathbb{C} \to \mathbb{C}$ so that $\varphi(X_i) = X_i'$ for all $1 \leq i \leq n$. In practice the collections of objects under consideration will not be labeled X_i and X_i', but there will be some natural bijection between them. Then our requirement is that φ respects this natural bijection. See Figures 5.1 and 5.2 for some examples.

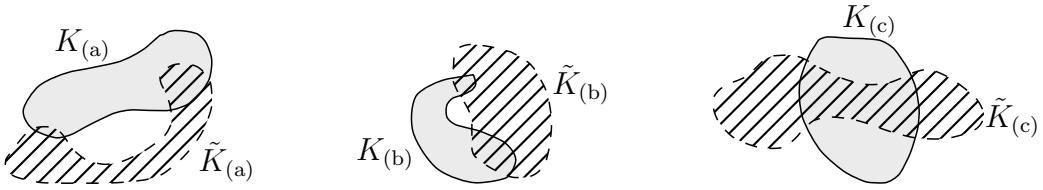

Figure 5.1: **Example topological configurations of pairs of closed Jordan domains.** The pair $\{K_{(a)}, \tilde{K}_{(a)}\}$ is in the same topological configuration as is the pair $\{K_{(b)}, \tilde{K}_{(b)}\}$. Here we use the natural associations $K_{(a)} \leftrightarrow K_{(b)}$ and $\tilde{K}_{(a)} \leftrightarrow \tilde{K}_{(b)}$. However, both pairs $\{K_{(a)}, \tilde{K}_{(a)}\}$ and $\{K_{(b)}, \tilde{K}_{(b)}\}$ are in different topological configurations from the pair $\{K_{(c)}, \tilde{K}_{(c)}\}$.

We will say that certain conditions on some objects *uniquely determine* their topological configuration if any two collections of objects satisfying the given conditions are in the same topological configuration.

Example 5.1. Suppose $C \subset \mathbb{C}$ is a circle, and $z \in \mathbb{C}$ is a point. Then the topological configuration of $\{C, z\}$ is uniquely determined by whether z lies in the open disk bounded by C, along ∂C, or in the complement of the closed disk bounded by C.

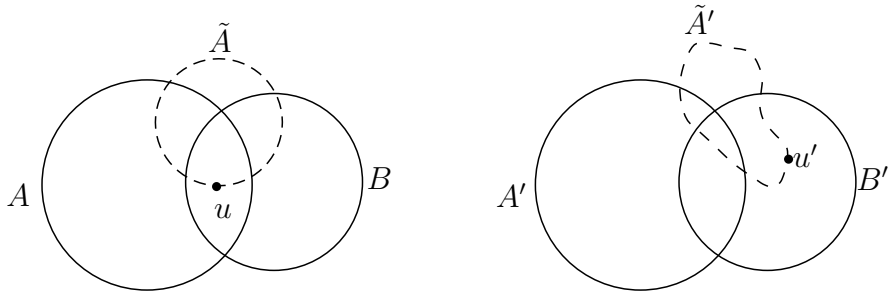

Figure 5.2: **More examples of topological configurations of sets in the plane.** The sets $A, B, \tilde{A}, A', B', \tilde{A}'$ are closed Jordan domains. We have that $\{A, B, \tilde{A}\}$ and $\{A', B', \tilde{A}'\}$ are in the same topological configuration, but $\{A, B, \tilde{A}, u\}$ and $\{A', B', \tilde{A}', u'\}$ are not. This is because $u \in A$, so any orientation-preserving homeomorphism $\varphi : \mathbb{C} \to \mathbb{C}$ sending A to A' must send u to a point inside of A', but $u' \notin A'$.

5.1 Invariance of fixed-point index under topological equivalence

The following lemma says that when working with fixed-point index, we need to consider our Jordan domains only "up to topological configuration."

Lemma 5.2. *Suppose K and $\tilde{K}$ are closed Jordan domains. Let $f : \partial K \to \partial \tilde{K}$ be an indexable homeomorphism. Suppose that K' and $\tilde{K}'$ are also closed Jordan domains, so that $\{K, \tilde{K}\}$ and $\{K', \tilde{K}'\}$ are in the same topological configuration, via the homeomorphism φ. Let $f' : \partial K' \to \partial \tilde{K}'$ be induced in the natural way, explicitly as $f' = \varphi|_{\partial \tilde{K}} \circ f \circ \varphi^{-1}|_{\partial K'}$. Then f' is indexable with respect to the usual orientation on $\partial K'$ and $\partial \tilde{K}'$, and $\eta(f) = \eta(f')$.*

Proof. The following is well-known. For a reference, see Chapters 1 and 2 of [FM12].

Fact 5.3. *Every orientation-preserving homeomorphism $\mathbb{C} \to \mathbb{C}$ is homotopic to the identity map via homeomorphisms.*

Thus let $H_t : \mathbb{C} \times [0, 1] \to \mathbb{C}$ be such a homotopy from the identity to φ. Explicitly, for fixed t we have that H_t is an orientation-preserving homeomorphism $\mathbb{C} \to \mathbb{C}$, with H_0 equal to the identity on $\mathbb{C}$ and $H_1 = \varphi$.

Let $K_t = H_t(K)$ and $\tilde{K}_t = H_t(\tilde{K})$. Then K_t and $\tilde{K}_t$ are closed Jordan domains, because H_t is a homeomorphism. Let $f_t : \partial K_t \to \partial \tilde{K}_t$ be induced in the natural way, explicitly as $H_t|_{\partial \tilde{K}} \circ f \circ H_t^{-1}|_{\partial K_t}$. Let $\gamma_t = \{f_t(z) - z\}_{z \in \partial K_t}$. Then tautologically $\eta(f)$ is the winding number of γ_0 around the origin, and $\eta(f')$ is the winding number of γ_1 around the origin.

Every γ_t is a closed curve because ∂K_t is a closed curve and f_t is continuous. Once we establish that no γ_t passes through the origin Lemma 5.2 will be proved because we have an induced homotopy from γ_0 to γ_1, and two curves homotopic in $\mathbb{C} \setminus \{0\}$ have the same winding number around the origin. Suppose for contradiction that $0 \in \gamma_t$. Then there is a $z \in \partial K_t$ so that $f_t(z) = z$. Thus $H_t \circ f \circ H_t^{-1}(z) = z$, and so $f(H_t^{-1}(z)) = H_t^{-1}(z)$, contradicting the fixed-point-free condition on f. □

5.2 General position for closed Jordan domains

Lemma 5.4. *Suppose K and $\tilde{K}$ are closed Jordan domains in general position. Then ∂K and $\partial \tilde{K}$ meet a finite number of times. Suppose that $z \in \partial K \cap \partial \tilde{K}$. Orient ∂K and $\partial \tilde{K}$ positively as usual. Then one of the following two mutually exclusive possibilities holds at the point z.*

1. *The curve $\partial \tilde{K}$ is entering K, and the curve ∂K is exiting $\tilde{K}$.*
2. *The curve ∂K is entering $\tilde{K}$, and the curve $\partial \tilde{K}$ is exiting K.*

Thus as we traverse ∂K, we alternate arriving at points of $\partial K \cap \partial \tilde{K}$ where (1) occurs and those where (2) occurs, and the same holds as we traverse $\partial \tilde{K}$. In particular, this implies that ∂K and $\partial \tilde{K}$ meet an even number of times.

Proof. The curves ∂K and $\partial \tilde{K}$ meet finitely often by compactness and the general position hypothesis. Let $z \in \partial K \cap \partial \tilde{K}$. We may assume, by applying a homeomorphism, that locally near z the picture looks like Figure 5.3, with ∂K oriented down-to-up as shown. Then K lies to the left of ∂K. Now, certainly $\partial \tilde{K}$ is either entering or exiting K at z. Suppose $\partial \tilde{K}$ is entering K at z. Then $\partial \tilde{K}$ is oriented right-to-left, and so $\tilde{K}$ is below $\partial \tilde{K}$. Thus ∂K is exiting $\tilde{K}$, and case (1) occurs. Similarly, if $\partial \tilde{K}$ is exiting K at z then ∂K is entering $\tilde{K}$ at z, so case (2) occurs. □

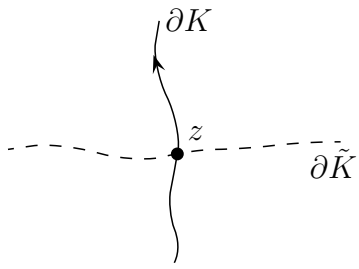

Figure 5.3: **A meeting point between two Jordan curves in general position.** The orientation shown on ∂K implies that K lies to the left. Depending on the orientation chosen for $\partial \tilde{K}$ we will get that $\tilde{K}$ lies above ∂K or below it.

5.3 Topological configurations of convex closed Jordan domains

Lemma 5.5. *Let K and $\tilde{K}$ be convex closed Jordan domains in general position, so that ∂K and $\partial\tilde{K}$ meet $2M > 0$ times. Suppose that K' and $\tilde{K}'$ are also convex closed Jordan domains in general position so that ∂K and $\partial\tilde{K}$ meet $2M > 0$ times. Then $\{K, \tilde{K}\}$ and $\{K', \tilde{K}'\}$ are in the same topological configuration.*

Remark 5.6. Figure 5.1 shows that Lemma 5.5 is not true if we omit the condition that the closed Jordan domains are convex. The clean construction we use in our proof is due to Nic Ford and Jordan Watkins. Refer to Figure 5.4 for an example for the construction.

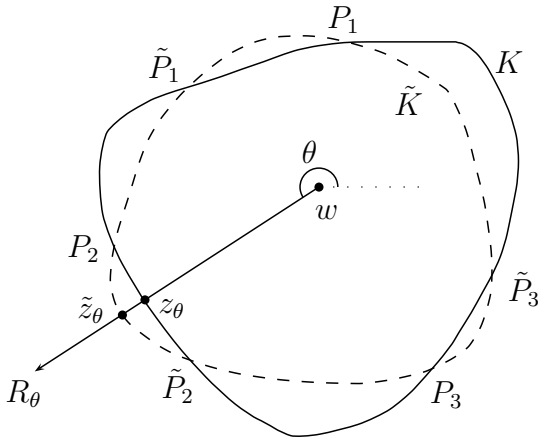

Figure 5.4: **Two convex closed Jordan domains K and $\tilde{K}$ in general position, with boundaries meeting at six points.** As θ varies positively, the ray R_θ scans around the boundaries of both K and $\tilde{K}$ positively.

Proof. Let w be a common interior point of K and $\tilde{K}$. Let R_θ be the ray emanating from the point w at an angle of θ from the positive real direction. Then for any θ we have that each of $R_\theta \cap \partial K$ and $R_\theta \cap \partial\tilde{K}$ consists of a single point. This is a consequence of the following fact, which is routine to prove.

Fact 5.7. *Suppose that the set $X \subset \mathbb{C}$ is convex. Suppose that S is a closed straight line segment contained in X. If exactly one endpoint of S lies in ∂X, then no other point of S lies in ∂X.*

Thus denote $\{z_\theta\} = R_\theta \cap \partial K$ and $\{\tilde{z}_\theta\} = R_\theta \cap \partial\tilde{K}$. Then as θ varies from 0 to 2π, we have that z_θ traverses ∂K in the positive direction and that $\tilde{z}_\theta$ traverses $\partial\tilde{K}$ in the positive direction.

Observation 5.8. *We have that $z_\theta \in \partial K \cap \partial \tilde{K}$ if and only if $\tilde{z}_\theta \in \partial K \cap \partial \tilde{K}$ if and only if $z_\theta = \tilde{z}_\theta$.*

Recall Lemma 5.4 and our hypothesis that ∂K and $\partial \tilde{K}$ meet at $2M > 0$ points. We denote by $P_1, \ldots, P_M \in \partial K \cap \partial \tilde{K}$ the points where ∂K is entering $\tilde{K}$, and by $\tilde{P}_1, \ldots, \tilde{P}_M \in \partial K \cap \partial \tilde{K}$ the points where $\partial \tilde{K}$ is entering K. See Figure 5.4 for examples of points P_i and $\tilde{P}_i$. We abuse notation and consider the indices of the P_i and $\tilde{P}_i$ only modulo M. For example, we write $P_{M+2} = P_2$. Then as θ varies positively the points z_θ and $\tilde{z}_\theta$ arrive at the P_i and the $\tilde{P}_i$ in the same cyclic order by Observation 5.8. Label so that as θ varies positively, we arrive at the P_i and the $\tilde{P}_i$ in the following cyclic order:

$$P_1, \tilde{P}_1, P_2, \tilde{P}_2, \ldots, P_M, \tilde{P}_M, P_1, \tilde{P}_1, \ldots$$

Suppose K' and $\tilde{K}'$ are as in the statement of Lemma 5.5. We wish to construct an orientation-preserving homeomorphism $\mathbb{C} \to \mathbb{C}$ sending K to K' and $\tilde{K}$ to $\tilde{K}'$. For every object under consideration defined for K and $\tilde{K}$, define the analogous object for K' and $\tilde{K}'$ the same way, and denote it with a prime. For example w' is a common interior point of K' and $\tilde{K}'$.

Let $\mathbb{S}^1$ denote $[0, 2\pi]$ with the endpoints identified. Let $\psi : \mathbb{S}^1 \to \mathbb{S}^1$ be a homeomorphism so that $z_\theta = \tilde{z}_\theta = P_i$ if and only if $z'_{\psi(\theta)} = \tilde{z}'_{\psi(\theta)} = P'_i$ and $z_\theta = \tilde{z}_\theta = \tilde{P}_i$ if and only if $z'_{\psi(\theta)} = \tilde{z}'_{\psi(\theta)} = \tilde{P}'_i$. We will use a prime to denote the image of θ under this homeomorphism, so $\psi(\theta) = \theta'$.

Observation 5.9. *For every θ there is a piecewise linear homeomorphism $\varphi_\theta : R_\theta \to R'_{\theta'}$ sending w to w', z_θ to $z'_{\theta'}$, and $\tilde{z}_\theta$ to $\tilde{z}'_{\theta'}$. We insist that φ_θ is an isometry on the unbounded component of $R_\theta \setminus \{z_\theta, \tilde{z}_\theta\}$.*

This is routine to check from our definitions. Then the φ_θ glue together to a homeomorphism $\mathbb{C} \to \mathbb{C}$ sending K to K' and $\tilde{K}$ to $\tilde{K}'$. This homeomorphism is orientation-preserving because it preserves the orientations $\partial K \to \partial K'$ and $\partial \tilde{K} \to \partial \tilde{K}'$ by construction. □

Remark 5.10. Which *a priori* topological configurations can occur for two Jordan curves in general position is a poorly understood question, and is known as the study of *meanders*. We are fortunate that our setting is nice enough that a statement like that of Lemma 5.5 is possible. Thanks to Thomas Lam for informing us of the topic of meander theory.

5.4 Eliminating irrelevant topological configurations

This chapter is devoted to the following proposition.

Proposition 5.11. *Suppose that* $\{A, B\}$ *and* $\{\tilde{A}, \tilde{B}\}$ *are pairs of overlapping closed disks in the plane* $\mathbb{C}$ *in general position. Suppose that* $A \setminus B$ *meets* $\tilde{A} \setminus \tilde{B}$*, that* $A \cap B$ *meets* $\tilde{A} \cap \tilde{B}$*, and that* $B \setminus A$ *meets* $\tilde{B} \setminus \tilde{A}$*. Then given any three of the disks* $A, B, \tilde{A}, \tilde{B}$*, the topological configuration of those three disks is one of those depicted in the appropriate one of figures* ♢*,* ♡*,* ♠*,* ♣*.*

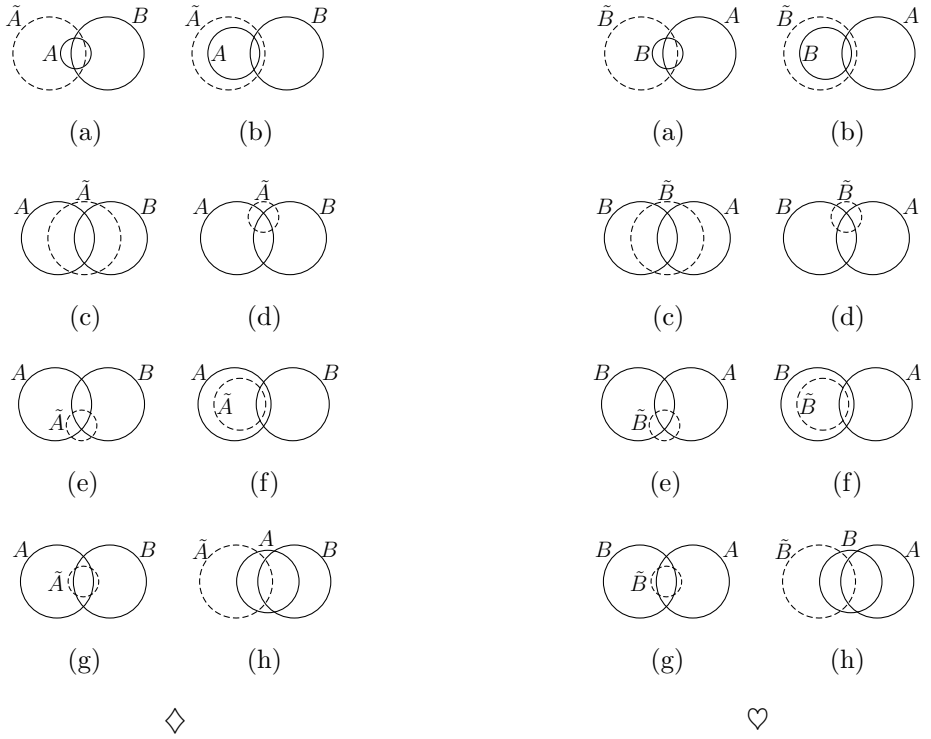

Figures ♢, ♡: **The possible topological configurations of** $\{A, B, \tilde{A}\}$ **and** $\{A, B, \tilde{B}\}$ **respectively, under the hypotheses of Proposition 5.11.**

We will often make reference to the configurations depicted in figures ♢, ♡, ♠, ♣, so for convenience we have placed another copy of them in an appendix, starting on p. 142. If the appropriate three-disk subset of $\{A, B, \tilde{A}, \tilde{B}\}$ is in a topological configuration depicted in the respective one of figures ♢, ♡, ♠, ♣, we will indicate this simply by saying that the corresponding configuration *occurs*, for example that ♢a occurs.

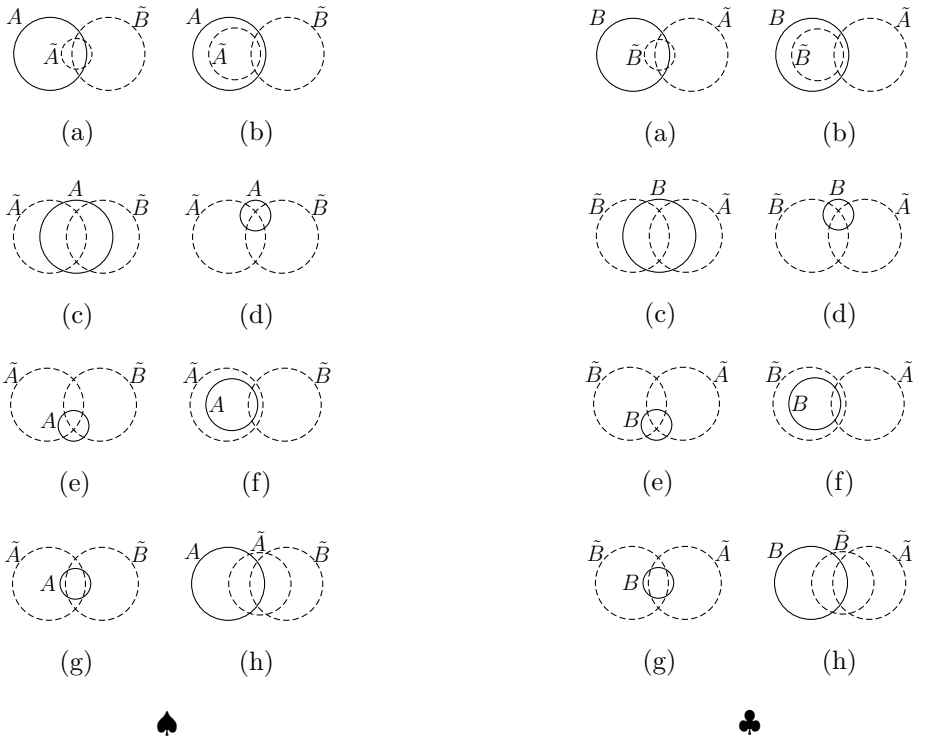

♠ ♣

Figures ♠, ♣: **The possible topological configurations of $\{A, \tilde{A}, \tilde{B}\}$ and $\{B, \tilde{A}, \tilde{B}\}$ respectively, under the hypotheses of Proposition 5.11.**

5.4.1 The topological configuration of an ellipse relative to the four quadrants

Lemma 5.12. *Suppose that S and S' are ellipses in $\mathbb{C}$ so that*

- *neither S nor S' passes through the origin,*
- *neither S nor S' is tangent to either of the real and imaginary axes, and*
- *each of S and S' passes through the same subset of the four quadrants.*

Then there is an orientation-preserving homeomorphism $\mathbb{C} \to \mathbb{C}$ sending S to S', fixing each of the four quadrants set-wise.

Proof. The proof breaks into two cases, one where both S and S' pass through all four quadrants, and the other where they do not.

Case 1. Suppose that S and S' pass through all four quadrants.

In this case, the domains bounded by S and S' both contain the origin as an interior point. Let R_θ be the ray emanating from the origin at an angle of θ from the positive

real axis. Denote $\{z_\theta\} = R_\theta \cap S$ and $\{z'_\theta\} = R_\theta \cap S'$. Let $\varphi_\theta : R_\theta \to R_\theta$ be defined piecewise as the natural linear identification $[0 \to z_\theta]_{R_\theta} \to [0 \to z'_\theta]_{R_\theta}$ and the natural isometry $[z_\theta \to \infty)_{R_\theta} \to [z'_\theta \to \infty)_{R_\theta}$. Then the φ_θ glue to a homeomorphism $\mathbb{C} \to \mathbb{C}$ that does what we want.

Case 2. Suppose that S and S' miss at least one quadrant.

Let R_θ be defined as before and let $\mathbb{S}^1$ denote $[0, 2\pi]$ with the endpoints identified. There are two values of θ for which R_θ is tangent to S. Denote them θ_1 and θ_2, labeled so that as θ varies positively from θ_1 to θ_2, the region swept out by R_θ contains S. Define θ'_1 and θ'_2 analogously.

Let $\psi : \mathbb{S}^1 \to \mathbb{S}^1$ be an orientation-preserving homeomorphism satisfying

- that $\psi(\theta) = \theta$ if $\theta = 0, \pi/4, \pi/2, 3\pi/4$, and
- that $\psi(\theta_1) = \theta'_1$ and $\psi(\theta_2) = \theta'_2$.

Such a ψ exists because S and S' meet the same subset of the four quadrants by hypothesis, so R_{θ_1} and $R_{\theta'_1}$ lie in the same quadrant, and R_{θ_2} and $R_{\theta'_2}$ lie in the same quadrant. There is a simple special case to consider, where R_{θ_1}, R_{θ_2}, $R_{\theta'_1}$, $R_{\theta'_2}$ all lie in the same quadrant, the details of which are left to the reader.

Let $\{p_1\} = S \cap R_{\theta_1}$ and $\{p_2\} = S \cap R_{\theta_2}$. Let K be a convex closed Jordan domain containing the origin as an interior point, satisfying $S \cap \partial K = \{p_1, p_2\}$. Let K' be chosen analogously. For every θ let $\psi_\theta : R_\theta \to R_{\psi(\theta)}$ be a homeomorphism which is piecewise linear in the natural way from the components of $R_\theta \setminus (S \cup \partial K)$ to those of $R_{\psi(\theta)} \setminus (S' \cup \partial K')$, so that ψ_θ is an isometry on the unbounded component of $R_\theta \setminus (S \cup \partial K)$. Then the ψ_θ glue to a homeomorphism $\mathbb{C} \to \mathbb{C}$ that does what we want. □

5.4.2 Reducing the proof of Proposition 5.11 to finitely many cases

Orient ∂A and ∂B positively as usual. Let $\{u, v\} = \partial A \cap \partial B$. Label u and v so that u is the point of $\partial A \cap \partial B$ where ∂A enters B, and v is the point of $\partial A \cap \partial B$ where ∂B enters A. See Figure 5.11 for an example.

Lemma 5.13. *Suppose that $\mathbb{C} \setminus (A \cup B \cup \tilde{A})$ is connected. Then the topological configuration of A, B, and $\tilde{A}$ is uniquely determined by which of the four components of $\mathbb{C} \setminus (\partial A \cup \partial B)$ the circle $\partial \tilde{A}$ passes through and whether $v \in \tilde{A}$.*

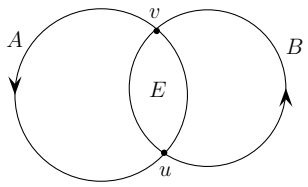

Figure 5.11: **The definitions of u and v in terms of the orientations on ∂A and ∂B.**

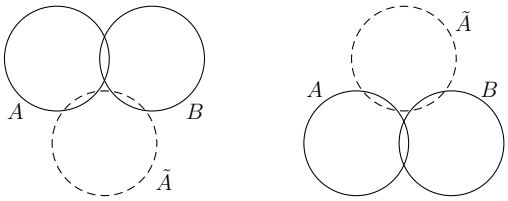

Figure 5.12: **Two different topological configurations of three disks $\{A, B, \tilde{A}\}$, where $\partial\tilde{A}$ passes through the same components of $\mathbb{C} \setminus (\partial A \cup \partial B)$ in both cases.** We see that $\tilde{A}$ and $A \cap B$ do not meet in either case, so this example should not worry us too much in light of the hypotheses of Proposition 5.11.

Remark 5.14. Figure 5.12 provides a counterexample to Lemma 5.13 if we remove the hypothesis that $\mathbb{C} \setminus (A \cup B \cup \tilde{A})$ is connected.

Proof. We construct the homeomorphism that does what we want by composition. First apply a Möbius transformation m sending v to ∞ and u to the origin. Then $m(\partial A)$ and $m(\partial B)$ are lines meeting transversely at the origin. By the general position hypothesis we have that $m(\partial\tilde{A})$ is a circle. Identifying $\mathbb{C}$ with $\mathbb{R}^2$, let $L : \mathbb{R}^2 \to \mathbb{R}^2$ be an orientation-preserving linear map sending $m(\partial A)$ to the imaginary axis, and $m(\partial B)$ to the real axis. Denote $\varphi = L \circ m$. It follows from the definition of u that then the right half-plane is $\varphi(A)$ and the lower half-plane is $\varphi(B)$. Thus the four open quadrants are the images under φ of the components of $\hat{\mathbb{C}} \setminus (\partial A \cup \partial B)$, and $\varphi(\partial\tilde{A}) =: E$ is an ellipse.

Let A', B', and $\tilde{A}'$ be closed disks in $\mathbb{C}$, satisfying hypotheses analogous to those on A, B, and $\tilde{A}$. For every object under consideration defined for A, B, and $\tilde{A}$, define the respective object for A', B', and $\tilde{A}'$ the same way, and denote it with a prime. Suppose

- that $\mathbb{C} \setminus (A' \cup B' \cup \tilde{A}')$ is connected,
- that $\partial\tilde{A}$ passes through a given component of $\mathbb{C} \setminus (\partial A \cup \partial B)$ if and only if $\partial\tilde{A}'$ passes through the respective component of $\mathbb{C} \setminus (\partial A' \cup \partial B')$, and
- that $v \in \tilde{A}$ if and only if $v' \in \tilde{A}'$.

Then $E' = \varphi'(\partial \tilde{A}')$ is an ellipse which passes through the same quadrants as does E. Let $\psi : \mathbb{C} \to \mathbb{C}$ be an orientation-preserving homeomorphism sending E to E', fixing every quadrant set-wise.

Now $\psi \circ \varphi(\hat{\mathbb{C}} \setminus (A \cup B \cup \tilde{A})) = \varphi'(\hat{\mathbb{C}} \setminus (A' \cup B' \cup \tilde{A}')) =: U$ because $v \in \tilde{A}$ if and only if $v' \in \tilde{A}'$. Furthermore U is a connected open set.

Fact 5.15. *Suppose $V \subset \mathbb{C}$ is a connected open set, and let $z_1, z_2 \in V$. Then there is a homeomorphism $V \to V$ sending z_1 to z_2 which limits to the identity on ∂V.*

Thus let $\vartheta : \mathbb{C} \to \mathbb{C}$ be a homeomorphism which is the identity outside of U, and which sends $\psi \circ \varphi(\infty)$ to $\varphi'(\infty)$. Then $(\varphi')^{-1} \circ \vartheta \circ \psi \circ \varphi$ is an orientation-preserving homeomorphism $\mathbb{C} \to \mathbb{C}$ sending A to A', B to B', and $\tilde{A}$ to $\tilde{A}'$, completing the proof of Lemma 5.13. □

The following lemma allows us to actually apply Lemma 5.13 in our proof of Proposition 5.11.

Lemma 5.16. *Suppose the hypotheses of Proposition 5.11 hold. Then the set $\hat{\mathbb{C}} \setminus (A \cup B \cup \tilde{A})$ is a topological open disk.*

Proof. Let φ be as in the proof of Lemma 5.13. Suppose that $v \in \tilde{A}$. Then $\varphi(\hat{\mathbb{C}} \setminus \tilde{A})$ is the open domain bounded by the ellipse $E = \varphi(\partial \tilde{A})$. This open domain is convex, and the open second quadrant, which is equal to $\varphi(\hat{\mathbb{C}} \setminus (A \cup B))$, is also convex. Therefore their intersection, which is equal to $\varphi(\hat{\mathbb{C}} \setminus (A \cup B \cup \tilde{A}))$, is a non-empty convex open set in $\mathbb{C}$, thus a topological open disk, so $\hat{\mathbb{C}} \setminus (A \cup B \cup \tilde{A})$ is a topological open disk.

Next, suppose that $v \notin \tilde{A}$. Then $\varphi(\tilde{A})$ is the closed Jordan domain bounded by the ellipse $\varphi(\partial \tilde{A})$. The hypotheses of Proposition 5.11 imply that $\varphi(\tilde{A})$ must meet the third quadrant $\varphi(A \cap B)$. The second quadrant is equal to $\varphi(\mathbb{C} \setminus (A \cup B))$, so if $\hat{\mathbb{C}} \setminus (A \cup B \cup \tilde{A})$ is not a topological open disk, then $\varphi(\tilde{A})$ must meet the second quadrant. But an ellipse meeting the second and fourth quadrants must meet all four quadrants, so $\varphi(\tilde{A})$ is the closed Jordan domain bounded by an ellipse around the origin. It follows that the open second quadrant minus $\varphi(\tilde{A})$ is a topological open disk, so $\varphi(\mathbb{C} \setminus (A \cup B \cup \tilde{A}))$ is a topological open disk. □

5.4.3 Proof of Proposition 5.11

We prove that the disks $\{A, B, \tilde{A}\}$ are in one of the topological configurations depicted in figure $\diamondsuit$, and the rest of Proposition 5.11 follows by symmetry. Lemma 5.13 allows

us to enumerate the possible topological configurations of the three disks $\{A, B, \tilde{A}\}$ under the hypothesis that $\hat{\mathbb{C}} \setminus (A \cup B \cup \tilde{A})$ is connected. By Lemma 5.16 if A, B, and $\tilde{A}$ satisfy the hypotheses of Proposition 5.11 then their topological configuration will be represented on this list. Note that by the general position hypothesis, the circle $\partial\tilde{A}$ passes through a given component U of $\mathbb{C} \setminus (\partial A \cup \partial B)$ if and only if it passes through the closure of U, or through the union of U and some subset of ∂U. Thus we enumerate the components of $\mathbb{C} \setminus (\partial A \cup \partial B)$ as $A \setminus B$, $A \cap B$, $B \setminus A$, and $\mathbb{C} \setminus (A \cup B)$. The following observation will be our source of contradictions to the hypotheses of Proposition 5.11.

Observation 5.17. *Suppose that the hypotheses of Proposition 5.11 hold. Then we have*

- *that $\tilde{A}$ meets both $A \setminus B$ and $A \cap B$, and*
- *that $B \setminus A$ is not contained in $\tilde{A}$.*

Example 5.18. Suppose that A, B, and $\tilde{A}$ satisfy the following requirements:

- that $\hat{\mathbb{C}} \setminus (A \cup B \cup \tilde{A})$ is connected,
- that $\partial\tilde{A}$ meets $A \setminus B$ and $A \cap B$,
- that $\partial\tilde{A}$ does not meet $B \setminus A$ nor $\mathbb{C} \setminus (A \cup B)$, and
- that $v \in \tilde{A}$.

The disks depicted in Figure 5.13a satisfy these requirements, so we conclude by Lemma 5.13 that the topological configuration of $\{A, B, \tilde{A}\}$ is the one shown in Figure 5.13a. From the figure we see that $\tilde{A}$ contains $B \setminus A$, so by Observation 5.17, the hypotheses of Proposition 5.11 will be violated. We conclude that there are no disks A, B, and $\tilde{A}$ simultaneously satisfying the requirements given in this example and the hypotheses of Proposition 5.11.

The following lemma restricts the remaining cases to check.

Lemma 5.19. *Suppose the hypotheses of Proposition 5.11 hold. Then the following hold.*

1. *The circle $\partial\tilde{A}$ passes through at least two components of $\mathbb{C} \setminus (\partial A \cup \partial B)$.*

2. *We call two of the four components of* $\mathbb{C} \setminus (\partial A \cup \partial B)$ adjacent *if their closures share a side, that is, if they meet along a positive-length arc. If* $\partial \tilde{A}$ *passes through two non-adjacent components, then it passes through at least one of the other two components, and possibly through both other components.*

3. *Suppose* $\partial \tilde{A}$ *does not meet* $\mathbb{C} \setminus (A \cup B)$*. Then* $\tilde{A}$ *does not contain* v*.*

4. *Suppose that* $\partial \tilde{A}$ *passes through* $A \setminus B$*,* $B \setminus A$*, and* $\mathbb{C} \setminus (A \cup B)$*, but not* $A \cap B$*. Then* $v \in \tilde{A}$*.*

Proof. Throughout the proof let φ be as in the proof of Lemma 5.13. We will make implicit reference to the following fact throughout the rest of the paper.

Fact 5.20. *Suppose that* K_1 *and* K_2 *are closed Jordan domains. If* $\partial K_1 \subset K_2$*, then* $K_1 \subset K_2$*. If* ∂K_1 *and* ∂K_2 *do not meet, then either* K_1 *and* K_2 *are disjoint, or one contains the other.*

(1) Suppose for contradiction that $\partial \tilde{A}$ meets only one of the four components. In each case we will deduce a contradiction to the hypotheses of Proposition 5.11 via Observation 5.17. If $\partial \tilde{A}$ is contained in $A \setminus B$, then so is $\tilde{A}$, so $\tilde{A}$ does not meet $A \cap B$. Similarly if $\partial \tilde{A}$ is contained in $A \cap B$ then $\tilde{A}$ does not meet $A \setminus B$, and if $\partial \tilde{A}$ is contained in $B \setminus A$ then $\tilde{A}$ does not meet $A \setminus B$. Finally suppose $\partial \tilde{A}$ is contained in $\mathbb{C} \setminus (A \cup B)$. Then either $\tilde{A}$ is contained in $\mathbb{C} \setminus (A \cup B)$, in which case $\tilde{A}$ does not meet $A \cap B$, or $A \cup B$ is contained in $\tilde{A}$, in which case $B \setminus A$ is contained in $\tilde{A}$.

(2) If $\partial \tilde{A}$ passes through two non-adjacent components but neither of the other two, then the ellipse $\varphi(\partial \tilde{A})$ passes through two opposite quadrants, but neither of the two remaining quadrants.

(3) Suppose that $\partial \tilde{A}$ is contained in the bounded set $A \cup B$. Then $\tilde{A} \subset A \cup B$, and by the general position hypothesis $\tilde{A}$ is contained in the interior of $A \cup B$. But $v \in \partial(A \cup B)$, so $v \notin \tilde{A}$.

(4) Suppose for contradiction that $v \notin \tilde{A}$. Then $\varphi(\tilde{A})$ is the closed Jordan domain bounded by the ellipse $\varphi(\partial \tilde{A}) = E$. The ellipse E meets the first, second, and third quadrants, but not the fourth quadrant. Thus $\varphi(\tilde{A})$ is itself contained in the first, second, and third quadrants. This implies that $\tilde{A}$ does not meet $A \cap B$, because $A \cap B$ is sent to the fourth quadrant by φ. This is a contradiction to the hypotheses of Proposition 5.11 by Observation 5.17. □

Figure 5.14 lists the remaining cases, completing the proof of Proposition 5.11. □

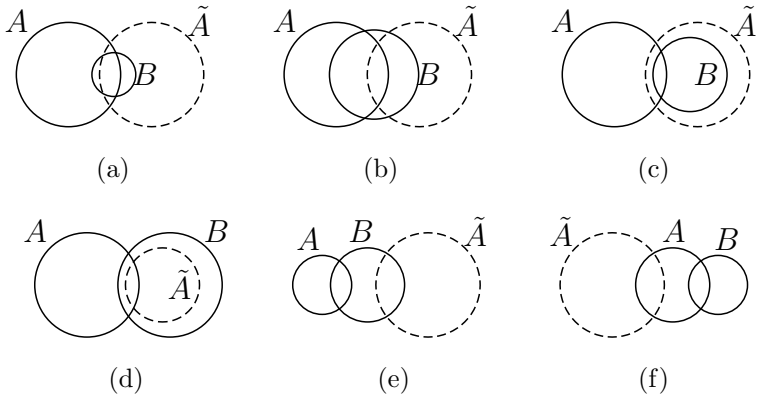

Figure 5.13: **Topological configurations for $\{A, B, \tilde{A}\}$ which guarantee a violation of the hypotheses of Proposition 5.11.** These are the remaining topological configurations that arise during our exhaustive proof of Proposition 5.11 via Figure 5.14.

If $\partial\tilde{A}$ meets				and	does Proposition 5.11 hold?	
$A \setminus B$,	$A \cap B$,	$B \setminus A$,	$\mathbb{C} \setminus (A \cup B)$	$v \in \tilde{A}$	If not, what fails?	(c.f. figure)
		✓	✓	✓	✓	♢b
✓			✓	✓	$(B \setminus A) \subset \tilde{A}$	5.13c
✓	✓				✓	♢f
	✓	✓			$(A \setminus B) \cap \tilde{A} = \varnothing$	5.13d
		✓	✓		$(A \setminus B) \cap \tilde{A} = \varnothing$	5.13e
✓			✓		$(A \cap B) \cap \tilde{A} = \varnothing$	5.13f
	✓	✓	✓	✓	✓	♢a
✓		✓	✓	✓	✓	♢c
✓	✓		✓	✓	$B \setminus A \subset \tilde{A}$	5.13a
✓	✓	✓			✓	♢g
	✓	✓	✓		$(A \setminus B) \cap \tilde{A} = \varnothing$	5.13b
✓	✓		✓		✓	♢h
✓	✓	✓	✓	✓	✓	♢d
✓	✓	✓	✓		✓	♢e

Figure 5.14: **The remaining cases to check for Proposition 5.11, after applying Lemma 5.19.**

5.5 Assorted lemmas on two pairs of disks

We take a moment to introduce some notation we use throughout the rest of the paper. Let γ be an oriented homeomorphic image of $\mathbb{S}^1$, for example a positively oriented Jordan curve. Let $a, b \in \gamma$ be distinct. Then $[a \to b]_\gamma$ is the oriented closed sub-arc of γ starting at a and ending at b. Then for example $[a \to b]_\gamma \cap [b \to a]_\gamma = \{a, b\}$ and $[a \to b]_\gamma \cup [b \to a]_\gamma = \gamma$.

Throughout this chapter let $\{A, B\}$ and $\{\tilde{A}, \tilde{B}\}$ be pairs of closed disks in general position, with boundaries oriented positively as usual. As before let $\{u, v\} = \partial A \cap \partial B$,

where we label u and v so that ∂A enters B at u and ∂B enters A at v. See Figure 5.11 for an example. We define and label $\tilde{u}$ and $\tilde{v}$ analogously. We denote $E = A \cap B$ and $\tilde{E} = \tilde{A} \cap \tilde{B}$.

This chapter is devoted to lemmas that we will need to handle some special cases later. The proofs are almost entirely self-contained. The reader is advised to skip this chapter for now, and refer back to the lemmas contained herein as they come up in later chapters.

Lemma 5.21. *The Jordan curves ∂E and $\partial \tilde{E}$ meet exactly 0, 2, 4, or 6 times.*

Proof. That they meet an even number of times is a consequence of the general position hypothesis. There is an immediate upper bound of 8 meeting points because each of ∂E and $\partial \tilde{E}$ is the union of two circular arcs. Suppose for contradiction that ∂E and $\partial \tilde{E}$ meet 8 times. Thus every meeting point of one of the circles ∂A and ∂B with one of $\partial \tilde{A}$ and $\partial \tilde{B}$ lies in $\partial E \cap \partial \tilde{E}$. It follows from the general position hypothesis that $\partial(A \cup B)$ does not meet $\partial \tilde{A}$ nor $\partial \tilde{B}$. But then one of three cases occurs: either $A \cup B$ is contained in one of the three bounded components of $\mathbb{C} \setminus (\partial \tilde{A} \cup \partial \tilde{B})$, or the sets $A \cup B$ and $\tilde{A} \cup \tilde{B}$, are disjoint, or the set $\tilde{A} \cup \tilde{B}$ is contained in $A \cup B$. If either of the first two of these cases occurs then it is impossible that ∂E and $\partial \tilde{E}$ meet at all, so we conclude that $\tilde{A} \cup \tilde{B} \subset A \cup B$. Then by symmetry we have $A \cup B \subset \tilde{A} \cup \tilde{B}$, so $A \cup B = \tilde{A} \cup \tilde{B}$, which forces a contradiction to the general position hypothesis. □

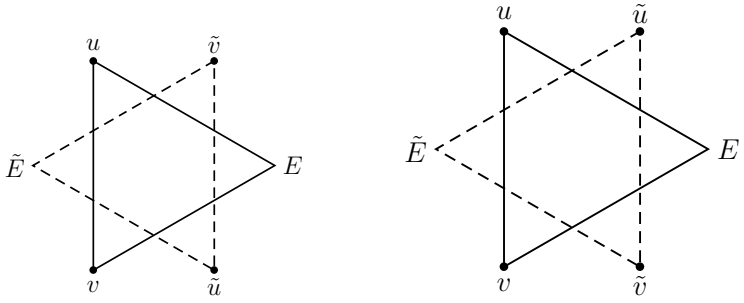

Figure 5.15: **The possible topological configurations for two eyes whose boundaries meet at six points.** Actually there are two remaining possibilities obtained by simultaneously swapping u with v and $\tilde{u}$ with $\tilde{v}$, but these are irrelevant for our applications because $\eta(f) = \eta(f^{-1})$ for indexable homeomorphisms f. (It might be nice to draw these with actual disks.)

Lemma 5.22. *Suppose that ∂E and $\partial \tilde{E}$ meet 6 times, that $A \setminus B$ and $\tilde{A} \setminus \tilde{B}$ meet, and that $B \setminus A$ and $\tilde{B} \setminus \tilde{A}$ meet. Then $\{E, u, v, \tilde{E}, \tilde{u}, \tilde{v}\}$ are in one of the two topological*

configurations represented in Figure 5.15, up to possibly simultaneously swapping u with v and $\tilde{u}$ with $\tilde{v}$.

Proof. By Lemma 5.5, if ∂E and $\partial \tilde{E}$ meet 6 times then they are in the topological configuration shown in Figure 5.16. We denote by ϵ_i the connected components of $\partial E \setminus \partial \tilde{E}$, and by $\tilde{\epsilon}_i$ the connected components of $\partial \tilde{E} \setminus \partial E$. Label the ϵ_i and $\tilde{\epsilon}_i$ as shown in Figure 5.16. We consider the indices of the ϵ_i and $\tilde{\epsilon}_i$ only modulo 6. For example, we write $\epsilon_{2+5} = \epsilon_1$.

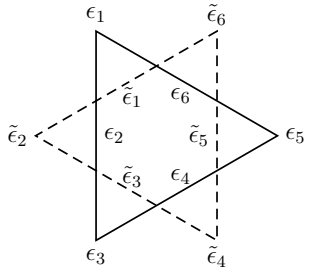

Figure 5.16: **The components of $\partial E \setminus \partial \tilde{E}$ and $\partial \tilde{E} \setminus \partial E$ for two convex closed Jordan domains E and $\tilde{E}$ meeting at six points.** The solid curve represents ∂E, and the dashed curve represents $\partial \tilde{E}$.

Proposition 5.11 allows us to make the following observation:

Observation 5.23. *Neither $\tilde{u}$ nor $\tilde{v}$ may lie in E, and neither u nor v may lie in $\tilde{E}$.*

To see why, note that if ∂E and $\partial \tilde{E}$ meet six times, then at least one of ∂A and ∂B must meet $\partial \tilde{E}$ at least three times. Thus at least one of ♠g and ♣g must occur. Thus $\tilde{u}$ and $\tilde{v}$ lie outside of at least one of A and B, but $E = A \cap B$, thus neither $\tilde{u}$ nor $\tilde{v}$ lies in E. The other part follows identically.

Thus we may assume that $u \in \epsilon_1$. Then v lies along ϵ_3 or ϵ_5. By relabeling the ϵ_i and switching the roles of u and v as necessary, we may assume that $v \in \epsilon_3$. Our proof will be done once we show that neither $\tilde{u}$ nor $\tilde{v}$ may lie along $\tilde{\epsilon}_2$. Suppose for contradiction that $\tilde{u}$ lies along $\tilde{\epsilon}_2$. Then $\tilde{v}$ lies along either $\tilde{\epsilon}_4$ or $\tilde{\epsilon}_6$. If $\tilde{v} \in \tilde{\epsilon}_4$, then the circular arc $[v \to u]_{\partial E}$ meets the circular arc $[\tilde{v} \to \tilde{u}]_{\partial \tilde{E}}$ three times, a contradiction. Similarly, if $\tilde{v} \in \tilde{\epsilon}_6$, then the circular arc $[v \to u]_{\partial E}$ meets the circular arc $[\tilde{u} \to \tilde{v}]_{\partial \tilde{E}}$ three times, also a contradiction. Thus $\tilde{u} \notin \tilde{\epsilon}_2$. The argument is the same if we had initially let $\tilde{v} \in \tilde{\epsilon}_2$. □

Lemma 5.24. *The following four statements hold.*

1. *If $[\tilde{u} \to \tilde{v}]_{\partial \tilde{E}}$ is contained in A and $[\tilde{v} \to \tilde{u}]_{\partial \tilde{E}}$ meets ∂A, then $B \setminus A$ and $\tilde{B} \setminus \tilde{A}$ are disjoint.*

2. *If* $[\tilde{v} \to \tilde{u}]_{\partial \tilde{E}}$ *is contained in* B *and* $[\tilde{u} \to \tilde{v}]_{\partial \tilde{E}}$ *meets* ∂B*, then* $A \setminus B$ *and* $\tilde{A} \setminus \tilde{B}$ *are disjoint.*

3. *If* $[u \to v]_{\partial E}$ *is contained in* $\tilde{A}$ *and* $[v \to u]_{\partial E}$ *meets* $\partial \tilde{A}$*, then* $\tilde{B} \setminus \tilde{A}$ *and* $B \setminus A$ *are disjoint.*

4. *If* $[v \to u]_{\partial E}$ *is contained in* $\tilde{B}$ *and* $[u \to v]_{\partial E}$ *meets* $\partial \tilde{B}$*, then* $\tilde{A} \setminus \tilde{B}$ *and* $A \setminus B$ *are disjoint.*

Proof. We prove only (1), as (2), (3), (4) are symmetric restatements of it. Suppose the hypotheses of (1) hold. Then both $\tilde{u}$ and $\tilde{v}$ lie in A. Thus the circular arc $[\tilde{v} \to \tilde{u}]_{\partial \tilde{E}}$ meets ∂A either exactly twice or not at all, in fact exactly twice because of the hypotheses. But $[\tilde{v} \to \tilde{u}]_{\partial \tilde{E}} = [\tilde{v} \to \tilde{u}]_{\partial \tilde{B}}$. Thus $[\tilde{u} \to \tilde{v}]_{\partial \tilde{B}}$ does not meet ∂A, and has its endpoints lying in A, so $[\tilde{u} \to \tilde{v}]_{\partial \tilde{B}} \subset A$.

From our definitions of $\tilde{u}$ and $\tilde{v}$, it is easy to check that $\partial(\tilde{B} \setminus \tilde{A})$ is the union of the arcs $[\tilde{u} \to \tilde{v}]_{\partial \tilde{B}}$ and $[\tilde{u} \to \tilde{v}]_{\partial \tilde{E}}$. It follows that $\partial(\tilde{B} \setminus \tilde{A})$ is contained in A. Thus $\tilde{B} \setminus \tilde{A}$ is contained in A, and so is disjoint from $B \setminus A$. □

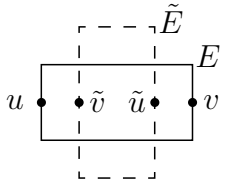

Figure 5.17: **A topological configuration of two eyes which guarantees that** $A \setminus B$ **and** $\tilde{A} \setminus \tilde{B}$ **do not meet, and that** $B \setminus A$ **and** $\tilde{B} \setminus \tilde{A}$ **do not meet.**

Lemma 5.25. *Suppose* $E, u, v, \tilde{E}, \tilde{u}, \tilde{v}$ *are in the topological configuration depicted in Figure 5.17. Then* $A \setminus B$ *and* $\tilde{A} \setminus \tilde{B}$ *do not meet, and* $B \setminus A$ *and* $\tilde{B} \setminus \tilde{A}$ *do not meet.*

Proof. The curves $\partial \tilde{A} \setminus \partial \tilde{E}$ and $\partial \tilde{B} \setminus \partial \tilde{E}$ both have $\tilde{u}$ and $\tilde{v}$ as their endpoints and otherwise avoid $\tilde{E}$. Thus each must cross ∂E twice. These four crossings together with the points $\partial E \cap \partial \tilde{E}$ accounts for all eight possible intersection points between $\partial A \cup \partial B$ and $\partial \tilde{A} \cup \partial \tilde{B}$. Thus the arc $[\tilde{v} \to \tilde{u}]_{\partial \tilde{E}}$ does not meet ∂B. Because this arc meets $B \supset E$, we conclude that $[\tilde{v} \to \tilde{u}]_{\partial \tilde{E}}$ is contained in B. Note that $[\tilde{u} \to \tilde{v}]_{\partial \tilde{E}}$ meets ∂B. Thus by part (1) of Lemma 5.25 we get that $A \setminus B$ and $\tilde{A} \setminus \tilde{B}$ are disjoint. That $B \setminus A$ and $\tilde{B} \setminus \tilde{A}$ are disjoint follows by symmetry. □

Lemma 5.26. *Suppose* $u \in \tilde{E}$ *and* $\tilde{u} \in E$*. Then* $A \setminus B$ *and* $\tilde{A} \setminus \tilde{B}$ *do not meet, or* $B \setminus A$ *and* $\tilde{B} \setminus \tilde{A}$ *do not meet.*

Proof. Suppose for contradiction that $u \in \tilde{E}$ and $\tilde{u} \in E$, but that $A \setminus B$ and $\tilde{A} \setminus \tilde{B}$ meet, and that $B \setminus A$ and $\tilde{B} \setminus \tilde{A}$ meet.

Observation 5.27. *Neither of $B \setminus A$ and $\tilde{B} \setminus \tilde{A}$ contains the other.*

To see why this is true, note that u is an interior point of $\tilde{E}$ by the general position hypothesis, and that $u \in \partial(B \setminus A)$. Thus $B \setminus A$ meets the exterior of $\tilde{B} \setminus \tilde{A}$. A similar argument gives that $\tilde{B} \setminus \tilde{A}$ meets the exterior of $B \setminus A$.

We are supposing for contradiction that $B \setminus A$ and $\tilde{B} \setminus \tilde{A}$ meet, and by Observation 5.27 neither of them contains the other. Thus if we can show that $\partial(B \setminus A)$ and $\partial(\tilde{B} \setminus \tilde{A})$ do not meet we will have derived a contradiction, as desired.

Note that Proposition 5.11 applies. This allows us to make the following observation.

Observation 5.28. *Either $\diamondsuit a$ or $\diamondsuit e$ occurs, either $\heartsuit a$ or $\heartsuit d$ occurs, either $\spadesuit a$ or $\spadesuit e$ occurs, and either $\clubsuit a$ or $\clubsuit d$ occurs.*

We prove that either $\diamondsuit$a or $\diamondsuit$e occurs, and the other parts of the observation follow similarly. Because $\tilde{E} \subset \tilde{A}$, we may eliminate any candidate topological configurations where $u \notin \tilde{A}$. This eliminates $\diamondsuit$d, $\diamondsuit$f, $\diamondsuit$g, and $\diamondsuit$h. Next, because $\tilde{u} \in \partial\tilde{A}$, we may eliminate any candidate topological configurations where $\partial\tilde{A}$ does not meet E, as this would preclude $\tilde{u} \in E$. This eliminates $\diamondsuit$b and $\diamondsuit$c, leaving us with only the two claimed possibilities. Thus the remainder of our proof breaks into cases as follows.

Case 1. Suppose that both $\diamondsuit$a and $\heartsuit$a occur.

Then $\partial(A \setminus B)$ is contained in $\tilde{A}$, and $\partial(B \setminus A)$ is contained in $\tilde{B}$. Thus $\partial(A \setminus B) \cup \partial(B \setminus A)$ is contained in $\tilde{A} \cup \tilde{B}$. But $\partial(A \cup B)$ is contained in $\partial(A \setminus B) \cup \partial(B \setminus A)$, thus in $\tilde{A} \cup \tilde{B}$. We conclude that $A \cup B \subset \tilde{A} \cup \tilde{B}$. Now $\tilde{u} \in \partial(\tilde{A} \cup \tilde{B})$ and $E \subset A \cup B$, so by the general position hypothesis we get a contradiction to $\tilde{u} \in E$.

Case 2. Suppose that both $\diamondsuit$a and $\heartsuit$d occur.

Then $u \in \tilde{E}$ and $v \in \tilde{A} \setminus \tilde{B}$. One of the following two sub-cases occurs.

Sub-case 2.1. Suppose that $\spadesuit$a occurs.

Then ∂A does not meet $\tilde{A} \setminus \tilde{B}$. But v lies on ∂A, contradicting $v \in \tilde{A} \setminus \tilde{B}$.

Sub-case 2.2. Suppose that $\spadesuit$e occurs. Then one of $\clubsuit$a and $\clubsuit$d occurs.

From $\spadesuit$e and that $u \in \tilde{E}$ and $v \in \tilde{A} \setminus \tilde{B}$, it follows that $\partial(B \setminus A) \cap \partial A = [u \to v]_{\partial E}$ does not meet $\partial(\tilde{B} \setminus \tilde{A})$. If $\clubsuit$a occurs, then $\partial B \supset \partial(B \setminus A) \cap \partial B$ does not meet $\partial(\tilde{B} \setminus \tilde{A})$. If $\clubsuit$d occurs, then via $u \in \tilde{E}$ and $v \in \tilde{A} \setminus \tilde{B}$ we get that $\partial(B \setminus A) \cap \partial B = [u \to v]_{\partial B}$

does not meet $\partial(\tilde{B} \setminus \tilde{A})$. In either case $\partial(B \setminus A)$ and $\partial(\tilde{B} \setminus \tilde{A})$ do not meet, giving us a contradiction.

Cases (1) and (2) together rule out $\heartsuit$a, $\spadesuit$a, and $\clubsuit$a by symmetry, so the only remaining case is the following.

Case 3. Suppose that $\diamondsuit$e, $\heartsuit$d, $\spadesuit$e, and $\clubsuit$d occur.

By $\diamondsuit$e and $\heartsuit$d we have that $u \in \tilde{E}$ and $v \in \mathbb{C} \setminus (\tilde{A} \cup \tilde{B})$. Then from $\spadesuit$e and $\clubsuit$d we get that neither $\partial(B \setminus A) \cap \partial A = [u \to v]_{\partial A}$ nor $\partial(B \setminus A) \cap \partial B = [u \to v]_{\partial B}$ meets $\partial(\tilde{B} \setminus \tilde{A})$, again giving us the desired contradiction. $\square$

Chapter 6

Torus parametrization

Let K and $\tilde{K}$ be closed Jordan domains in general position, so that ∂K and $\partial \tilde{K}$ meet at $2M \geq 0$ points, with boundaries oriented as usual. Let $\partial K \cap \partial \tilde{K} = \{P_1, \ldots, P_M, \tilde{P}_1, \ldots, \tilde{P}_M\}$, where P_i and $\tilde{P}_i$ are labeled so that at every P_i we have that ∂K is entering $\tilde{K}$, and at every $\tilde{P}_i$ we have that $\partial \tilde{K}$ is entering K. Imbue $\mathbb{S}^1$ with an orientation and let $\kappa : \partial K \to \mathbb{S}^1$ and $\tilde{\kappa} : \partial \tilde{K} \to \mathbb{S}^1$ be orientation-preserving homeomorphisms. We refer to this as fixing a *torus parametrization* for K and $\tilde{K}$.

We consider a point (x, y) on the 2-torus $\mathbb{T} = \mathbb{S}^1 \times \mathbb{S}^1$ to be parametrizing simultaneously a point $\kappa^{-1}(x) \in \partial K$ and a point $\tilde{\kappa}^{-1}(y) \in \partial \tilde{K}$. We denote by $p_i \in \mathbb{T}$ be the unique point $(x, y) \in \mathbb{T}$ satisfying $\kappa^{-1}(x) = \tilde{\kappa}^{-1}(y) = P_i$, similarly $\tilde{p}_i \in \mathbb{T}$. Note that by the general position hypothesis no pair of points in $\{p_1, \ldots, p_M, \tilde{p}_1, \ldots, \tilde{p}_M\}$ share a first coordinate, nor a second coordinate.

Suppose we pick $(x_0, y_0) \in \mathbb{S}^1 \times \mathbb{S}^1$. Then we may draw an image of $\mathbb{T} = \mathbb{S}^1 \times \mathbb{S}^1$ by letting $\{x_0\} \times \mathbb{S}^1$ be the vertical axis and letting $\mathbb{S}^1 \times \{y_0\}$ be the horizontal axis. Then we call (x_0, y_0) a *base point* for the drawing. See Figure 6.1 for an example.

Suppose that $f : \partial K \to \partial \tilde{K}$ is an orientation-preserving homeomorphism. Then f determines an oriented curve γ in $\mathbb{T}$ for us, namely its graph $\gamma = \{(\kappa(z), \tilde{\kappa}(f(z))\}_{z \in \partial K}$, with orientation obtained by traversing ∂K positively. Note that f is fixed-point-free if and only if its associated curve γ misses all of the p_i and $\tilde{p}_i$. Pick $u \in \partial K$ and denote $\tilde{u} = f(u)$. Then if we draw the torus parametrization for K and $\tilde{K}$ using the base point $(\kappa(u), \tilde{\kappa}(\tilde{u}))$, the curve γ associated to f "looks like the graph of a strictly increasing function." The converse is also true: given any such γ, it determines for us an orientation-preserving homeomorphism $\partial K \to \partial \tilde{K}$ sending u to $\tilde{u}$, which is fixed-point-free if and only if γ misses all of the p_i and $\tilde{p}_i$.

Before moving on to proofs, we would like to thank Jordan Watkins for fruitful discussions in on the content in this part of the thesis.

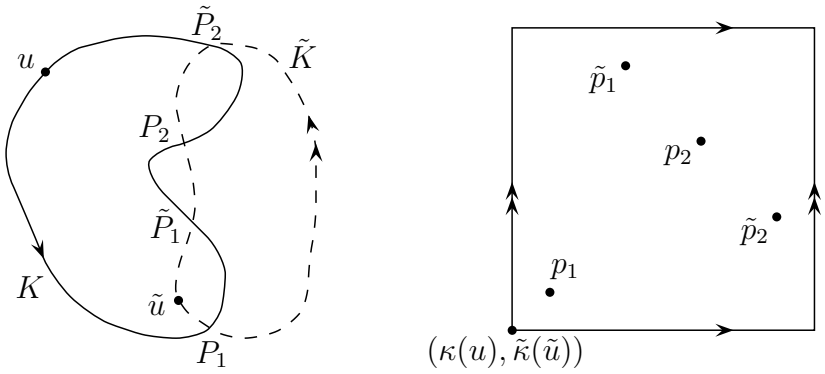

Figure 6.1: **A pair of closed Jordan domains K and $\tilde{K}$ and a torus parametrization for them, drawn with base point $(\kappa(u), \tilde{\kappa}(\tilde{u}))$.** The key points to check are that as we vary the first coordinate of $\mathbb{T}$ positively starting at u, we arrive at $\kappa(P_1)$, $\kappa(\tilde{P}_1)$, $\kappa(P_2)$, and $\kappa(\tilde{P}_2)$ in that order, and as we vary the second coordinate of $\mathbb{T}$ positively starting at $\tilde{\kappa}(\tilde{u})$, we arrive at $\tilde{\kappa}(P_1)$, $\tilde{\kappa}(\tilde{P}_2)$, $\tilde{\kappa}(P_2)$, and $\tilde{\kappa}(P_1)$ in that order.

6.1 Computing fixed-point index using a torus parametrization

Suppose that $f(u) = \tilde{u}$, equivalently that $(\kappa(u), \tilde{\kappa}(\tilde{u})) \in \gamma$. The curve γ and the horizontal and vertical axes $\{\tilde{\kappa}(\tilde{u})\} \times \mathbb{S}^1$ and $\mathbb{S}^1 \times \{\kappa(u)\}$ divide $\mathbb{T}$ into two simply connected open sets $\Delta_\uparrow(u, \gamma)$ and $\Delta_\downarrow(u, \gamma)$ as shown in Figure 6.2. We suppress the dependence on $\tilde{u}$ in the notation because $\tilde{u} = f(u)$. If neither $u \in \partial\tilde{K}$ nor $\tilde{u} \in \partial K$ then every p_i and every $\tilde{p}_i$ lies in either $\Delta_\downarrow(u, \gamma)$ or $\Delta_\uparrow(u, \gamma)$. In this case we write $\#p_\downarrow(u, \gamma)$ to denote $|\{p_1, \ldots, p_M\} \cap \Delta_\downarrow(u, \gamma)|$ the number of points p_i which lie in $\Delta_\downarrow(u, \gamma)$, and we define $\#p_\uparrow(u, \gamma)$, $\#\tilde{p}_\downarrow(u, \gamma)$, and $\#\tilde{p}_\uparrow(u, \gamma)$ in the analogous way. Denote by $\omega(\alpha, z)$ the winding number of the closed curve $\alpha \subset \mathbb{C}$ around the point $z \notin \alpha$.

Lemma 6.1. *Let K and $\tilde{K}$ be closed Jordan domains. Fix a torus parametrization of K and $\tilde{K}$ via κ and $\tilde{\kappa}$. Let $f : \partial K \to \partial\tilde{K}$ be an orientation-preserving fixed-point-free homeomorphism, with graph γ in $\mathbb{T}$. Suppose that $f(u) = \tilde{u}$, where $u \notin \partial\tilde{K}$ and $\tilde{u} \notin \partial K$. Then:*

$$\eta(f) = w(\gamma) = \omega(\partial K, \tilde{u}) + \omega(\partial\tilde{K}, u) - \#p_\downarrow(u, \gamma) + \#\tilde{p}_\downarrow(u, \gamma) \tag{6.1}$$

$$= \omega(\partial K, \tilde{u}) + \omega(\partial\tilde{K}, u) + \#p_\uparrow(u, \gamma) - \#\tilde{p}_\uparrow(u, \gamma) \tag{6.2}$$

The remainder of the section is spent proving Lemma 6.1.

Suppose γ_0 is any oriented closed curve in $\mathbb{T} \setminus \{p_1, \ldots, p_M, \tilde{p}_1, \ldots, \tilde{p}_M\}$. Then the closed curve $\{\tilde{\kappa}^{-1}(y) - \kappa^{-1}(x)\}_{(x,y)\in\gamma_0}$ misses the origin, and has a natural orientation

obtained by traversing γ_0 positively. We denote by $w(\gamma_0)$ the winding number around the origin of $\{\tilde{\kappa}^{-1}(y) - \kappa^{-1}(x)\}_{(x,y)\in\gamma_0}$.

Observation 6.2. *If γ_1 and γ_2 are homotopic in $\mathbb{T} \setminus \{p_1, \ldots, p_M, \tilde{p}_1, \ldots, \tilde{p}_M\}$ then $w(\gamma_1) = w(\gamma_2)$.*

This is because the homotopy between γ_1 and γ_2 in $\mathbb{T} \setminus \{p_2, \ldots, p_M, \tilde{p}_1, \ldots, \tilde{p}_M\}$ induces a homotopy between the closed curves $\{\tilde{\kappa}^{-1}(y) - \kappa^{-1}(x)\}_{(x,y)\in\gamma_1}$ and $\{\tilde{\kappa}^{-1}(y) - \kappa^{-1}(x)\}_{(x,y)\in\gamma_2}$ in the punctured plane $\mathbb{C} \setminus \{0\}$.

Suppose that $f : \partial K \to \partial\tilde{K}$ is a fixed-point-free orientation-preserving homeomorphism. Let γ be its graph in $\mathbb{T}$. If γ has orientation induced by traversing ∂K and $\partial\tilde{K}$ positively, then the following is a tautology.

Observation 6.3. $\eta(f) = w(\gamma)$

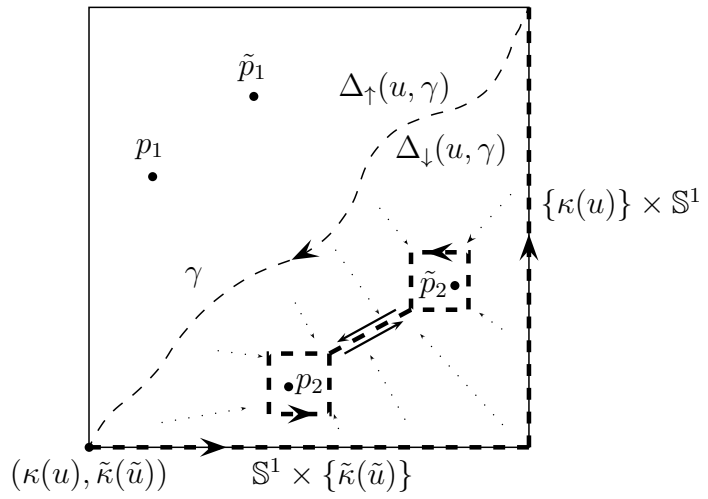

Figure 6.2: **A homotopy from $\partial\Delta_\downarrow(u,\gamma)$ to Γ.** Here the orientation shown on γ is the opposite of the orientation induced by traversing ∂K positively.

Orient $\partial\Delta_\downarrow(u,\gamma)$ as shown in Figure 6.2. Then $\partial\Delta_\downarrow(u,\gamma)$ is the concatenation of the curve γ traversed backwards with $\mathbb{S}^1 \times \{\tilde{\kappa}(\tilde{u})\}$ and $\{\kappa(u)\} \times \mathbb{S}^1$, where the two latter curves are oriented according to the positive orientation on $\mathbb{S}^1$.

Observation 6.4. *If $\mathbb{S}^1 \times \{\tilde{\kappa}(\tilde{u})\}$ and $\{\kappa(u)\} \times \mathbb{S}^1$ are oriented according to the positive orientation on $\mathbb{S}^1$, then $w(\mathbb{S}^1 \times \{\tilde{\kappa}(\tilde{u})\}) = \omega(\partial K, \tilde{u})$ and $w(\{\kappa(u)\} \times \mathbb{S}^1) = \omega(\partial\tilde{K}, u)$.*

It is also easy to see that if we concatenate two closed curves γ_1 and γ_2 that meet at a point, we get $w(\gamma_1 \circ \gamma_2) = w(\gamma_1) + w(\gamma_2)$. Thus in light of the orientations on

$\partial\Delta_\downarrow(u,\gamma)$ and all other curves concerned we get:

$$\begin{aligned} w(\partial\Delta_\downarrow(u,\gamma)) &= w(\mathbb{S}^1 \times \{\tilde{\kappa}(\tilde{u})\}) + w(\{\kappa(u)\} \times \mathbb{S}^1) - w(\gamma) \\ &= \omega(\partial K, \tilde{u}) + \omega(\partial \tilde{K}, u) - \eta(f) \end{aligned}$$

For every i let $\zeta(p_i)$ and $\zeta(\tilde{p}_i)$ be small squares around p_i and $\tilde{p}_i$ respectively in $\mathbb{T}$, oriented as shown in Figure 6.2. By *square* we mean a simple closed curve which decomposes into four "sides," so that on a given side one of the two coordinates of $\mathbb{S}^1 \times \mathbb{S}^1 = \mathbb{T}$ is constant. Pick the rectangles small enough so that the closed boxes they bound are pairwise disjoint and do not meet $\partial\Delta_\downarrow(u,\gamma)$.

Let Γ be the closed curve in $\Delta_\downarrow(u,\gamma)$ obtained in the following way. First, start with every loop $\zeta(p_i)$ and $\zeta(\tilde{p}_i)$ for those p_i and $\tilde{p}_i$ lying in $\Delta_\downarrow(u,\gamma)$. Let δ_0 be an arc contained in the interior of $\Delta_\downarrow(u,\gamma)$ which meets each $\zeta(p_i)$ and $\zeta(\tilde{p}_i)$ contained in $\Delta_\downarrow(u,\gamma)$ at exactly one point. It is easy to prove inductively that such an arc exists. Let δ be the closed curve obtained by traversing δ_0 first in one direction, then in the other. Then let Γ be obtained by concatenating δ with every $\zeta(p_i)$ and $\zeta(\tilde{p}_i)$ contained in $\Delta_\downarrow(u,\gamma)$.

Observation 6.5. *The curves Γ and $\partial\Delta_\downarrow(u,\gamma)$ are homotopic in $\mathbb{T} \setminus \{p_1, \ldots, p_M, \tilde{p}_1, \ldots, \tilde{p}_M\}$. Also $w(\delta) = 0$. It follows that:*

$$w(\partial\Delta_\downarrow(u,\gamma)) = w(\Gamma) = \sum_{p_i \in \Delta_\downarrow(u,\gamma)} w(\zeta(p_i)) + \sum_{\tilde{p}_i \in \Delta_\downarrow(u,\gamma)} w(\zeta(\tilde{p}_i))$$

See Figure 6.2 for an example. On the other hand, the following holds.

Observation 6.6. $w(\zeta(p_i)) = 1$, $w(\zeta(\tilde{p}_i)) = -1$

To see why, suppose that $\zeta(p_i) = \partial([x_0 \to x_1]_{\mathbb{S}^1} \times [y_0 \to y_1]_{\mathbb{S}^1})$. Then up to orientation-preserving homeomorphism the picture near P_i is as in Figure 6.3. We let (x,y) traverse $\zeta(p_i)$ positively starting at (x_0,y_0), keeping track of the vector $\tilde{\kappa}^{-1}(y) - \tilde{\kappa}^{-1}(x)$ as we do so. The vector $\tilde{\kappa}^{-1}(y_0) - \tilde{\kappa}^{-1}(x_0)$ points to the right. As x varies from x_0 to x_1, the vector $\tilde{\kappa}^{-1}(y) - \tilde{\kappa}^{-1}(x)$ rotates in the positive direction, that is, counter-clockwise, until it arrives at $\tilde{\kappa}^{-1}(y_0) - \kappa^{-1}(x_1)$, which points upward. Continuing in this fashion, we see that $\tilde{\kappa}^{-1}(y) - \tilde{\kappa}^{-1}(x)$ makes one full counter-clockwise rotation as we traverse $\zeta(p_i)$. The proof that $w(\zeta(\tilde{p}_i)) = -1$ is similar. Combining all of our observations establishes equation 6.1. The proof that equation 6.2 holds is similar. $\square$

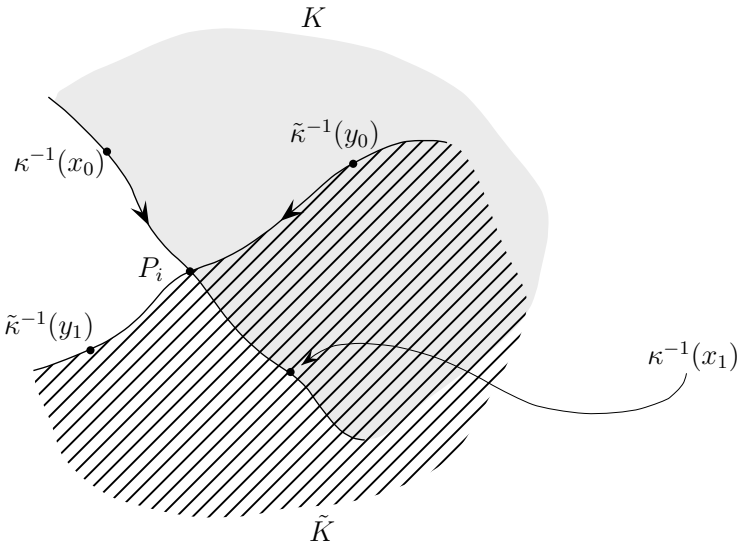

Figure 6.3: **The local picture near P_i.** This allows us to compute the "local fixed-point index" $w(\zeta(p_i))$ of f near P_i.

6.2 Proof of the Circle Index Lemma 3.1

We prove the Circle Index Lemma 3.1 as our first application of Lemma 6.1. We restate the claim of each part of the Circle Index Lemma 3.1 just before we prove it for the reader's convenience. The original statement is on p. 63. Let K and $\tilde{K}$ be closed Jordan domains, and let $f : \partial K \to \partial \tilde{K}$ be an indexable homeomorphism. We suppose for simplicity that ∂K and $\partial \tilde{K}$ are in general position, as we may assume this hypothesis for all of our applications. We remark that even when ∂K and $\partial \tilde{K}$ fail to be in general position, we may often suppose that they are anyway without loss of generality by applying Lemma 3.3. Suppose that we have fixed a torus parametrization for K and $\tilde{K}$ via κ and $\tilde{\kappa}$, and let γ be the graph of f as usual.

(1) We wish to show that $\eta(f) = \eta(f^{-1})$. Considering f^{-1} rather than f amounts to interchanging the roles of K and $\tilde{K}$, thus of κ and $\tilde{\kappa}$, etc. In particular this reverses the roles of $\Delta_\downarrow$ and $\Delta_\uparrow$, and makes every p_i into a $\tilde{p}_i$ and vice versa. It follows that $\eta(f) = \eta(f^{-1})$.

(2) We wish to show that if one of K and $\tilde{K}$ contains the other, then $\eta(f) = 1$. By part (1) we may suppose without loss of generality that $\tilde{K} \subset K$. Then there are no p_i and no $\tilde{p}_i$, and however u and $\tilde{u}$ are chosen we have that $\omega(\partial K, \tilde{u}) = 1$ and $\omega(\partial \tilde{K}, u) = 0$.

(3) We wish to show that if K and $\tilde{K}$ are disjoint, then $\eta(f) = 0$. Then there are no p_i and no $\tilde{p}_i$, as in the proof of part (2). This time though, however u and $\tilde{u}$ are chosen, both $\omega(\partial K, \tilde{u})$ and $\omega(\partial \tilde{K}, u)$ are 0.

(4) We wish to show that if ∂K and $\partial\tilde{K}$ meet at exactly two points, then $\eta(f) \geq 0$. We may pick u in the interior of $\tilde{K}$ so that $\tilde{u} = f(u) \notin \partial K$. Then $\omega(\partial\tilde{K}, u) = 1$. Also there is exactly one p_i and exactly one $\tilde{p}_i$, so $\#p_\downarrow(u,\gamma) - \#\tilde{p}_\downarrow(u,\gamma) \geq -1$. It follows that $\eta(f) \geq 0$. □

6.3 The situation if no eye is contained in its partner

Proposition 6.7. *Let $\{A, B\}$ and $\{\tilde{A}, \tilde{B}\}$ be pairs of overlapping closed disks in the plane $\mathbb{C}$ in general position. Suppose that neither of $E = A \cap B$ and $\tilde{E} = \tilde{A} \cap \tilde{B}$ contains the other. Suppose further that $A \setminus B$ and $\tilde{A} \setminus \tilde{B}$ meet, and that $B \setminus A$ and $\tilde{B} \setminus \tilde{A}$ meet. Then there is a faithful indexable homeomorphism $e : \partial E \to \partial\tilde{E}$ satisfying $\eta(e) = 0$.*

The rest of this chapter is spent proving Proposition 6.7. If ∂E and $\partial\tilde{E}$ do not meet, we get that E and $\tilde{E}$ are disjoint. Then any indexable homeomorphism $e : \partial E \to \partial\tilde{E}$ satisfies $\eta(e) = 0$. Thus suppose that ∂E and $\partial\tilde{E}$ meet. Fix a torus parametrization for E and $\tilde{E}$ via $\kappa : \partial E \to \mathbb{S}^1$ and $\tilde{\kappa} : \partial\tilde{E} \to \mathbb{S}^1$. As before denote by p_i the points of $\partial E \cap \partial\tilde{E}$ where ∂E is entering $\tilde{E}$, and by $\tilde{p}_i$ those where $\partial\tilde{E}$ is entering E. Note that ∂E and $\partial\tilde{E}$ meet at exactly 2, 4, or 6 points by Lemma 5.21. The proof breaks into these three cases.

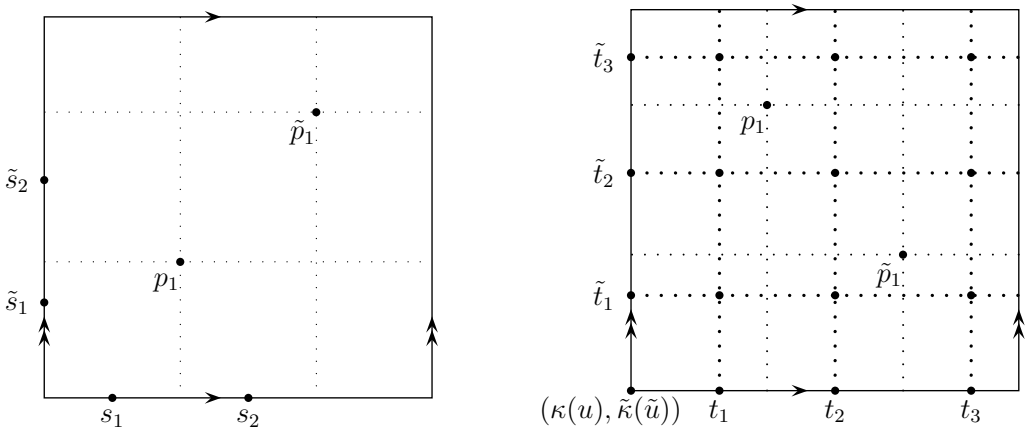

(a) (b) $(\kappa(u), \tilde{\kappa}(u)) = (s_1, \tilde{s}_2)$

Figure 6.4: **Two drawings of a torus parametrization for two eyes whose boundaries meet exactly twice.** There is some choice of base point giving the drawing on the left. The drawing on the right is the same torus parametrization drawn using the base point $(\kappa(u), \tilde{\kappa}(u)) = (s_1, \tilde{s}_2)$.

Case 1. Suppose that ∂E and $\partial\tilde{E}$ meet at exactly two points.

Then with an appropriate choice of base point, the torus parametrization for E and $\tilde{E}$ is as shown in Figure 6.4a. The points $s_1, s_2 \in \mathbb{S}^1$ in Figure 6.4a are exactly the topologically distinct places where $\kappa(u)$ may be, similarly $\tilde{s}_1, \tilde{s}_2 \in \mathbb{S}^1$ for $\tilde{\kappa}(\tilde{u})$. A choice of $(s_j, \tilde{s}_{\tilde{j}}) = (\kappa(u), \tilde{\kappa}(\tilde{u}))$ completely determines the topological configuration of $\{E, \tilde{E}, u, \tilde{u}\}$, and conversely every possible topological configuration of those sets is achieved via this procedure. By Lemma 5.26 we may suppose without loss of generality that $u \notin \tilde{E}$, thus that $\kappa(u) = s_1$.

Suppose first that $\tilde{\kappa}(\tilde{u}) = \tilde{s}_2$. In Figure 6.4b we redraw the torus parametrization for E and $\tilde{E}$ using the base point $(s_1, \tilde{s}_2) = (\kappa(u), \tilde{\kappa}(\tilde{u}))$. Then the points $t_1, t_2, t_3 \in \mathbb{S}^1$ are exactly the topologically distinct places where $\kappa(v)$ may be, similarly $\tilde{t}_1, \tilde{t}_2, \tilde{t}_3 \in \mathbb{S}^1$ for $\tilde{\kappa}(\tilde{v})$.

Observation 6.8. *A choice of $(s_j, \tilde{s}_{\tilde{j}}) = (\kappa(u), \tilde{\kappa}(\tilde{u}))$ and a subsequent choice of $(t_k, \tilde{t}_{\tilde{k}}) = (\kappa(v), \tilde{\kappa}(\tilde{v}))$ together completely determine the topological configuration of $\{E, \tilde{E}, u, \tilde{u}, v, \tilde{v}\}$. Conversely every possible topological configuration of $\{E, \tilde{E}, u, \tilde{u}, v, \tilde{v}\}$ is achieved by some choice of $(s_j, \tilde{s}_{\tilde{j}})$, and then a subsequent choice of $(t_k, \tilde{t}_{\tilde{k}})$, for $(\kappa(u), \tilde{\kappa}(\tilde{u}))$ and $(\kappa(v), \tilde{\kappa}(\tilde{v}))$ respectively.*

We are currently working under the assumption that $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_1, \tilde{s}_2)$. For every choice of $(t_k, \tilde{t}_{\tilde{k}}) = (\kappa(v), \tilde{\kappa}(\tilde{v}))$ we hope to find a faithful indexable homeomorphism $e : \partial E \to \partial \tilde{E}$ so that $\eta(e) = 0$.

Observation 6.9. *Suppose we have drawn the parametrization for E and $\tilde{E}$ using $(s_j, \tilde{s}_{\tilde{j}}) = (\kappa(u), \tilde{\kappa}(\tilde{u}))$ as the base point. Then finding a faithful indexable homeomorphism $e : \partial E \to \partial \tilde{E}$ amounts to finding a curve γ in $\mathbb{T} \backslash \{p_1, \ldots, \tilde{p}_M, \tilde{p}_1, \ldots, \tilde{p}_M\}$ which "looks like the graph of a strictly increasing function," from the lower-left-hand corner $(\kappa(u), \tilde{\kappa}(u))$ to the upper-right-hand corner, passing through $(\kappa(v), \tilde{\kappa}(\tilde{v})) = (t_k, \tilde{t}_{\tilde{k}})$. Having fixed such a curve γ, we may compute $\eta(e)$, where e is the homeomorphism associated to γ, using Lemma 6.1.*

In our current situation $\kappa(u) = s_1$ implies that $u \notin \tilde{K}$, and $\tilde{\kappa}(\tilde{u}) = \tilde{s}_2$ implies that $\tilde{u} \notin K$. Thus by Lemma 6.1 we wish to find curves γ so that both p_1 and $\tilde{p}_1$ lie in the upper diagonal $\Delta_{\uparrow}(u, \gamma)$, or both lie in the lower diagonal $\Delta_{\downarrow}(u, \gamma)$. Figure 6.5b depicts such a γ for every $(t_k, \tilde{t}_{\tilde{k}})$ except for $(t_2, \tilde{t}_2)$. Suppose $(t_2, \tilde{t}_2) = (\kappa(v), \tilde{\kappa}(\tilde{v}))$. Then $v \in \tilde{K}$ and $\tilde{v} \in K$, so we get a contradiction by Lemma 5.26. From now on points $(t_k, \tilde{t}_{\tilde{k}})$ which are handled via Lemma 5.26 will be labeled with an asterisk, as in Figure 6.5b.

Next suppose that $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_1, \tilde{s}_1)$. The situation is depicted in Figure 6.5a. Then $u \notin \tilde{K}$ and $\tilde{u} \in K$, so to achieve $\eta(e) = 0$ we wish to find curves γ so that

$p_1 \in \Delta_\downarrow(u,\gamma)$ and $\tilde{p}_1 \in \Delta_\uparrow(u,\gamma)$. This time there are four $(t_k, \tilde{t}_{\tilde{k}})$ for which this is not possible. For $(\kappa(v), \tilde{\kappa}(\tilde{v}) = (t_2, \tilde{t}_1), (t_2, \tilde{t}_3)$ we again get contradictions via Lemma 5.26. The following observation will be helpful for $(\kappa(v), \tilde{\kappa}(\tilde{v}) = (t_1, \tilde{t}_3), (t_3, \tilde{t}_1)$.

Observation 6.10. *Choose $(s_j, \tilde{s}_{\tilde{j}}) = (\kappa(u), \tilde{\kappa}(\tilde{u}))$ and draw our torus parametrization for E and $\tilde{E}$ using $(\kappa(u), \tilde{\kappa}(\tilde{u}))$ as the base point. Then a choice of $(t_k, \tilde{t}_{\tilde{k}}) = (\kappa(v), \tilde{\kappa}(\tilde{v}))$ defines for us four "quadrants," namely $[\kappa(u) \to \kappa(v)]_{\mathbb{S}^1} \times [\tilde{\kappa}(\tilde{u}) \to \tilde{\kappa}(\tilde{v})]_{\mathbb{S}^1}$ the points "below and to the left of" $(t_k, \tilde{t}_{\tilde{k}})$, etc. Then which of the two arcs $\partial A \cap \partial E$ and $\partial B \cap \partial E$, and which of $\partial \tilde{A} \cap \partial \tilde{E}$ and $\partial \tilde{B} \cap \partial \tilde{E}$, a point P_i or $\tilde{P}_i$ lies on is determined by which quadrant p_i or $\tilde{p}_i$ lies in.*

For example, suppose $(\kappa(v), \tilde{\kappa}(\tilde{v}) = (t_1, \tilde{t}_3)$. Then p_1 and $\tilde{p}_1$ lie in the lower-right-hand quadrant $[\kappa(v) \to \kappa(u)]_{\mathbb{S}^1} \times [\tilde{\kappa}(\tilde{u}) \to \tilde{\kappa}(\tilde{v})]_{\mathbb{S}^1}$, so both P_1 and $\tilde{P}_1$ lie on $\partial E \cap \partial B = [v \to u]_{\partial E}$ and on $\partial \tilde{E} \cap \partial \tilde{A} = [\tilde{u} \to \tilde{v}]_{\partial \tilde{E}}$. Also $[\tilde{v} \to \tilde{u}]_{\partial E}$ is contained in E, because both $\tilde{v}$ and $\tilde{u}$ are, and no p_i nor any $\tilde{p}_i$ lies the two upper quadrants $[\tilde{v} \to \tilde{u}]_{\mathbb{S}^1} \times \mathbb{S}^1$. Then we get a contradiction via Lemma 5.24. A similar argument gives us a contradiction via Lemma 5.24 for $(\kappa(v), \tilde{\kappa}(\tilde{v}) = (t_3, \tilde{t}_1)$. From now on points $(t_k, \tilde{t}_{\tilde{k}})$ which are handled via Lemma 5.24 in this way will be labeled with a diamond, as in Figure 6.5a. This completes the proof of Proposition 6.7 when ∂E and $\partial \tilde{E}$ meet at exactly two points.

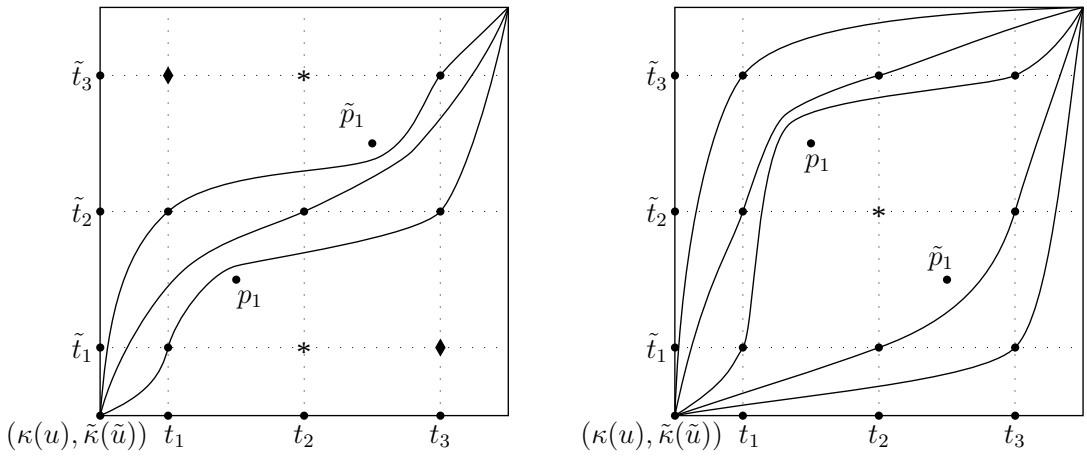

(a) $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_1, \tilde{s}_1)$ (b) $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_1, \tilde{s}_2)$

Figure 6.5: **Graphs of homeomorphisms e giving $\eta(e) = 0$ for a pair of eyes whose boundaries meet twice.** The torus parametrizations are drawn using the indicated choice of base point.

Case 2. Suppose that ∂E and $\partial \tilde{E}$ meet at exactly four points.

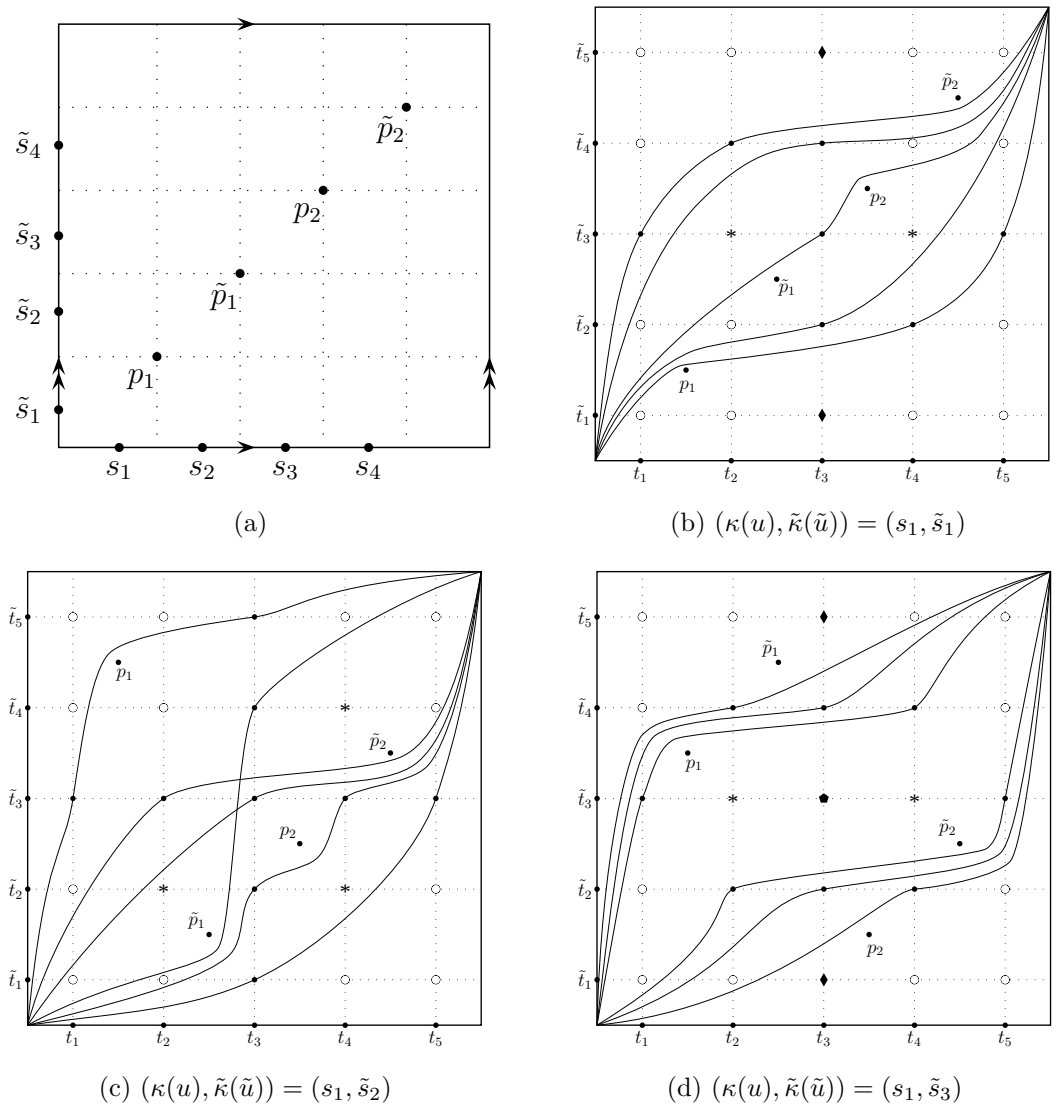

(a)

(b) $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_1, \tilde{s}_1)$

(c) $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_1, \tilde{s}_2)$

(d) $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_1, \tilde{s}_3)$

Figure 6.6: **The situation if two eyes' boundaries meet four times.** Figure (a) shows the torus parametrization for E and $\tilde{E}$ with some suitable choice of base point. Figures (b)–(d) give graphs of homeomorphisms e giving $\eta(e) = 0$, with torus parametrizations drawn using base point $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_j, \tilde{s}_{\tilde{j}})$ as indicated.

Lemma 5.5 guarantees that with a correct choice of base point, the torus parametrization for E and $\tilde{E}$ is as in Figure 6.6a. As before, we may suppose without loss of generality that $u \notin \tilde{K}$, thus $\kappa(u) = s_1$, by Lemma 5.26 and relabeling the s_i if necessary. Thus we have the possibilities $\tilde{\kappa}(\tilde{u}) = \tilde{s}_1, \tilde{s}_2, \tilde{s}_3, \tilde{s}_4$ to consider. The cases $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_1, \tilde{s}_2)$ and $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_1, \tilde{s}_4)$ are symmetric by Figure 6.7. Figures 6.6b–6.6d give the solutions for $\tilde{\kappa}(u) = \tilde{s}_1, \tilde{s}_2, \tilde{s}_3$, modulo some remaining special cases.

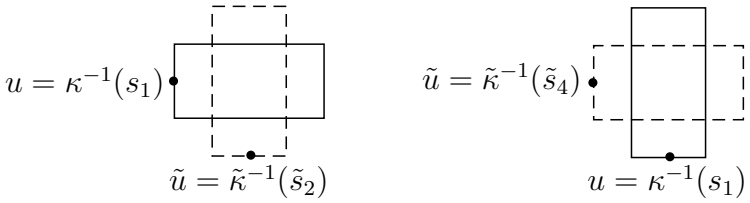

Figure 6.7: **The topological configurations of $\{E, u, \tilde{u}\}$ leading to the cases $(\kappa(u), \tilde{\kappa}(u)) = (s_1, \tilde{s}_2), (s_1, \tilde{s}_4)$.** We see that these are equivalent via a rotation, because $\eta(e) = \eta(e^{-1})$.

Points $(t_k, \tilde{t}_{\tilde{k}})$ labeled with an asterisk or a diamond are handled via Lemma 5.26 or 5.24 respectively as before. Suppose $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_1, \tilde{s}_1)$, and $(\kappa(v), \tilde{\kappa}(\tilde{v})) = (t_1, \tilde{t}_1)$ in Figure 6.6b. Then the upper-right-hand quadrant defined for us by $(t_1, \tilde{t}_1)$ contains all four points $p_1, p_2, \tilde{p}_1, \tilde{p}_2$, thus the circular arcs $[v \to u]_{\partial E}$ and $[\tilde{v} \to \tilde{u}]_{\partial \tilde{E}}$ meet four times, a contradiction. All points that are handled in this way are labeled with a small circle. Finally, if $(\kappa(u), \tilde{\kappa}(\tilde{u})) = (s_1, \tilde{s}_3)$ and $(\kappa(v), \tilde{\kappa}(\tilde{v})) = (t_3, \tilde{t}_3)$ in Figure 6.6d, we get a contradiction via Lemma 5.25.

Case 3. Suppose that ∂E and $\partial \tilde{E}$ meet at exactly six points.

Then Lemma 5.22 restricts us to two cases to consider. These are handled in Figure 6.8. This completes the proof of Proposition 6.7. □

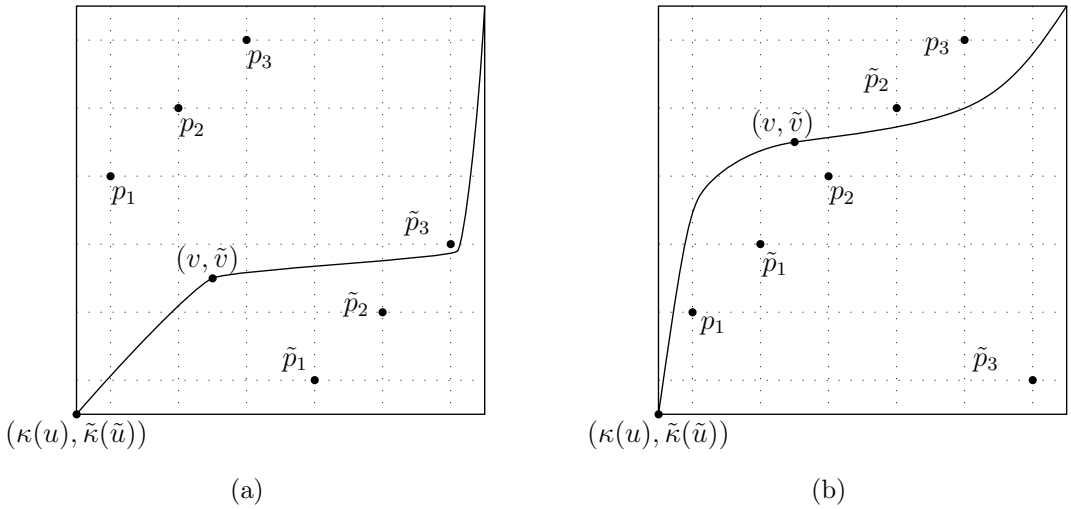

Figure 6.8: **Torus parametrizations for the eyes depicted in Figure 5.15.** Both are drawn with base point $(\kappa(u), \tilde{\kappa}(\tilde{u}))$. Each of the two curves is the graph of a faithful indexable homeomorphism $e : \partial E \to \partial \tilde{E}$ satisfying $\eta(e) = 0$.

6.4 The situation if no disk is contained in its partner

Proposition 6.11. *Let $\mathcal{D} = \{D_1, \ldots, D_n\}$ and $\tilde{\mathcal{D}} = \{\tilde{D}_1, \ldots, \tilde{D}_n\}$ be as in the statement of Theorem 3.8. That is, they are thin disk configurations in the plane $\mathbb{C}$ in general position, realizing the same pair (G, Θ) where $G = (V, E)$ is a graph and $\Theta : E \to [0, \pi)$. In addition, suppose that for all i, j the sets $D_i \setminus D_j$ and $\tilde{D}_i \setminus \tilde{D}_j$ meet. Suppose that there is no i so that one of D_i and $\tilde{D}_i$ contains the other. Suppose that for every disjoint non-empty $I, J \subset \{1, \ldots, n\}$ so that $I \sqcup J = \{1, \ldots, n\}$, there exists an eye E_{ij} with $i \in I$ and $j \in J$ so that one of E_{ij} and $\tilde{E}_{ij}$ contains the other. Then for every i we have that any faithful indexable homeomorphism $g_i : \partial D_i \to \partial \tilde{D}_i$ satisfies $\eta(g_i) \geq 1$. Furthermore there is a k so that D_i and D_k overlap for all i, and so that one of E_{ij} and $\tilde{E}_{ij}$ contains the other if and only if either $i = k$ or $j = k$.*

Recalling notation from before, if D_i and D_j overlap then $E_{ij} = D_i \cap D_j$, similarly $\tilde{E}_{ij}$, and a homeomorphism $g_i : \partial D_i \to \partial \tilde{D}_i$ is called *faithful* if it restricts to homeomorphisms $D_j \cap \partial D_i \to \tilde{D}_j \cap \partial D_i$ for all j. The rest of this chapter is spent proving Proposition 6.11.

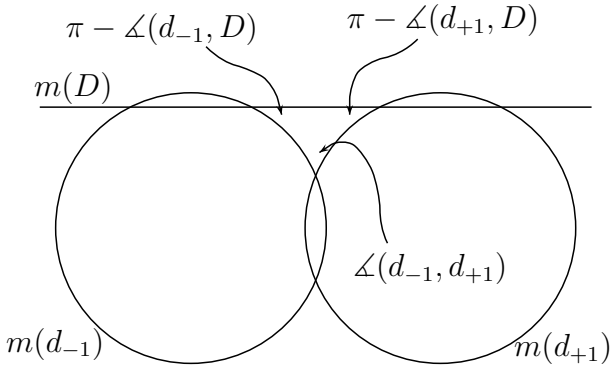

Figure 6.9: **The image of the Möbius transformation described in the proof of Lemma 6.12.**

Lemma 6.12. *Suppose that D, d_{-1}, d_{+1} are closed disks in the plane $\mathbb{C}$ in the topological configuration depicted in Figure 6.11a. Then $\pi + \measuredangle(d_{-1}, d_{+1}) < \measuredangle(d_{-1}, D) + \measuredangle(d_{+1}, D)$.*

Proof. Let m be a Möbius transformation sending a point on the bottom arc of $\partial D \setminus d_{-1} \cup d_{+1}$ to ∞, so that $m(D)$ is the lower half plane. Then the images of the disks under m are as depicted in Figure 6.9. We see that $(\pi - \measuredangle(d_{-1}, D)) + (\pi - \measuredangle(d_{+1}, D)) + \measuredangle(d_{-1}, d_{+1}) < \pi$ and the desired inequality follows. □

Lemma 6.13. *Suppose that D, d_{-1}, d_{+1} are closed disks in the plane $\mathbb{C}$ in the topological configuration depicted in Figure 6.10. Then $\measuredangle(d_{-1}, D) + \measuredangle(d_{+1}, D) < \pi + \measuredangle(d_{-1}, d_{+1})$.*

Proof. This is proved similarly to Lemma 6.12, see Figure 6.10. □

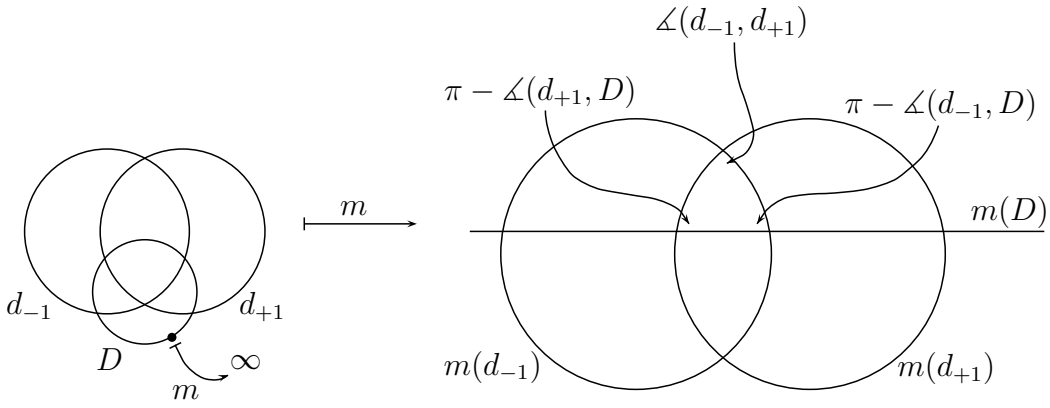

Figure 6.10: **A Möbius transformation chosen to prove Lemma 6.13.** Here $\partial m(D) = \mathbb{R}$, and $m(D)$ is the lower half-plane.

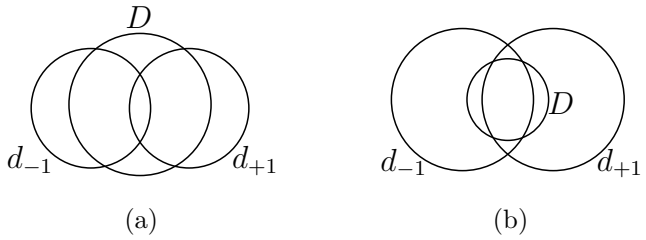

Figure 6.11: **The topological configurations for which we prove Lemma 6.14.**

Lemma 6.14. *Suppose that D, d_{-1}, d_{+1} are closed disks in the plane $\mathbb{C}$ in one of the two topological configurations depicted in Figure 6.11. In either case, we get that both $\measuredangle(d_{-1}, D)$ and $\measuredangle(d_{+1}, D)$ are strictly greater than $\measuredangle(d_{-1}, d_{+1})$.*

Proof. Suppose that the disks are in the configuration depicted in Figure 6.11a. Let m be a Möbius transformation sending a point on $\partial d_{+1} \setminus D$ to ∞. We may suppose without loss of generality that $m(d_{+1})$ is the lower half-plane. Then the image of our disks under m is as in Figure 6.12, where $\theta_1 = \measuredangle(d_{+1}, D)$ and $\theta_2 = \measuredangle(d_{-1}, d_{+1})$. It is then an easy exercise to show that $\theta_2 < \theta_1$ because the two circles $\partial m(d_{-1})$ and $\partial m(D)$ meet in the upper half-plane. The other inequality follows by symmetry. The case where the disks are in the configuration depicted in Figure 6.11b follows from

the first case after applying a Möbius transformation sending a point in the interior of $D \cap d_{-1} \cap d_{+1}$ to ∞. □

Claim 6.15. *Let i, j be so that $\tilde{E}_{ij} \subset E_{ij}$. Denote $A = D_i$, $B = D_j$, $\tilde{A} = \tilde{D}_i$, $\tilde{B} = \tilde{D}_j$. Then both ♠c and ♣c occur. Also one of ♢d, ♢e, ♢g occurs, and one of ♡d, ♡e, ♡g occurs. Furthermore at least one of ♢g and ♡g occurs.*

Proof. Both ♠c and ♣c occur, because these are the only candidates in Figure ♠, ♣ where $\tilde{A} \cap \tilde{B}$ is contained in the respective one of A and B. Note the following by Lemmas 6.14

$$\measuredangle(A, B) = \measuredangle(\tilde{A}, \tilde{B}) < \measuredangle(\tilde{A}, B) \tag{6.3}$$

and the following by Lemma 6.12.

$$\pi + \measuredangle(\tilde{A}, \tilde{B}) < \measuredangle(A, \tilde{A}) + \measuredangle(\tilde{B}, A), \qquad \pi + \measuredangle(\tilde{A}, \tilde{B}) < \measuredangle(\tilde{A}, B) + \measuredangle(\tilde{B}, B) \tag{6.4}$$

Next, because $\tilde{A} \cap \tilde{B}$ contains part of $\partial\tilde{A}$ and part of $\partial\tilde{B}$, both of these circles must pass through $A \cap B$. Noting that ♢f cannot occur because $\tilde{A} \not\subset A$, we conclude that one of ♢a, ♢d, ♢e, ♢g, and ♢h occurs. If either of ♢a and ♢h occurs, then Lemma 6.14 implies that $\measuredangle(\tilde{A}, B) < \measuredangle(A, B)$, contradicting 6.3. This leaves us with only the claimed possibilities ♢d, ♢e, and ♢g. By symmetry we also get that one of ♡d, ♡e, and ♡g occurs.

Finally, note by Lemma 6.13 that if ♢d or ♢e occurs then we get $\measuredangle(\tilde{A}, A) + \measuredangle(\tilde{A}, B) < \pi + \measuredangle(A, B)$, and if ♢d or ♢e occurs then we get $\measuredangle(\tilde{B}, A) + \measuredangle(\tilde{B}, B) < \pi + \measuredangle(A, B)$. We get that if neither of ♢g and ♡g occurs, then we may combine these two inequalities with 6.4 to arrive at a contradiction. □

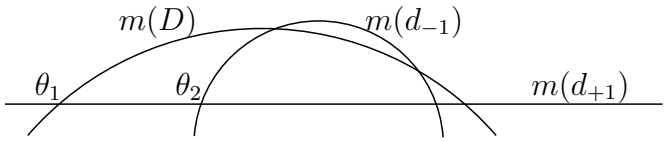

Figure 6.12: **The image of the Möbius transformation described in the proof of Lemma 6.14.**

Pick $1 \leq i \leq n$. By the hypotheses of Proposition 6.11 there is a j so that one of E_{ij} and $\tilde{E}_{ij}$ contains the other, without loss of generality so that $\tilde{E}_{ij} \subset E_{ij}$. Let $g_i : \partial D_i \to \partial \tilde{D}_i$ be a faithful indexable homeomorphism. Continuing with the notation of Lemma 6.15, regardless of which of ♢d, ♢e, and ♢g occurs, there is a point $z \in \partial A \cap \partial E$ so that z lies in the interior of $\tilde{A}$. Furthermore note that

$g_i(z) \in \partial\tilde{E}$ by the faithfulness condition, and that $\tilde{E} \subset A$ by our hypotheses, so $g_i(z)$ lies in the interior of A. Thus if we draw a torus parametrization for A and $\tilde{A}$ using $(\kappa(z), \tilde{\kappa}(g_i(z)))$ as the base point, Lemma 6.1 implies that $\eta(g_i) \geq 1$, because ∂A and $\partial\tilde{A}$ meet exactly twice. This establishes the first part of Proposition 6.11.

Next, let G_{u} be the undirected simple graph having $\{1, \ldots, n\}$ as its vertex set, so that $\langle i, j\rangle$ is an edge in G_{u} if and only if D_i and D_j overlap and one of E_{ij} and $\tilde{E}_{ij}$ contains the other. Note that G_{u} is connected, otherwise we could pick I to be the vertex set of one connected component of G_{u} and J to be $\{1, \ldots, n\} \setminus I$ to contradict the hypotheses of Proposition 6.11.

Let G be the directed graph obtained from G_{u} in the following way. Suppose $\langle i, j\rangle$ is an edge in G_{u}. Denote $A = D_i$, $B = D_j$, $\tilde{A} = \tilde{D}_i$, $\tilde{B} = \tilde{D}_j$. Then $\langle i \to j\rangle$ is an edge in G if and only if one of $\diamondsuit$g and $\spadesuit$g occurs. In particular Lemma 6.15 implies that if $\langle i, j\rangle$ is an edge in G_{u} then at least one of $\langle i \to j\rangle$ and $\langle j \to i\rangle$ is an edge in G, and possibly both are.

Claim 6.16. *Suppose that $\langle i \to j\rangle$ is an edge in G. Then $\langle i, j\rangle$ is the only edge in G_{u} having i as a vertex.*

Proof. Note that if $\diamondsuit$d or $\diamondsuit$e occurs then one intersection point $\partial A \cap \partial\tilde{A}$ lie in the interior of B, and if $\diamondsuit$g occurs then both do. Suppose without loss of generality that $\tilde{D}_i \cap \tilde{D}_j \subset D_i \cap D_j$. Then both intersection points $\partial D_i \cap \partial\tilde{D}_i$ lie in the interior of D_j. For contradiction let $k \neq j$ so that $\langle i, k\rangle$ is an edge in G_{u}. There are two cases.

Case 1. Suppose that $\tilde{D}_i \cap \tilde{D}_k \subset D_i \cap D_k$.

Then one or both points $\partial D_i \cap \partial\tilde{D}_i$ lie in the interior of D_k. Then there is a point in the interior of D_i which lies in the interiors of both D_j and D_k, a contradiction.

Case 2. Suppose that $D_i \cap D_k \subset \tilde{D}_i \cap \tilde{D}_k$.

Then by a symmetric restatement of Lemma 6.15 we get that both points $\partial D_i \cap \partial D_k$ lie in $\tilde{D}_i$. On the other hand $\tilde{D}_i \cap \partial D_i$ is contained in the interior of D_j by $\diamondsuit$a. Thus there are points interior to all of D_i, D_j, D_k, a contradiction. □

Thus G is either the graph on two vertices $\{i, j\}$ having one or both of $\langle i \to j\rangle$ and $\langle j \to i\rangle$ as edges, or is a graph having $\{k, i_1, \ldots, i_{n-1}\}$ as vertices and exactly the edges $\langle i_\ell \to k\rangle$ for $1 \leq \ell < n$. The last part of Proposition 6.11 follows. □

6.5 Proof of the Three Point Prescription Lemma 3.5

We restate the lemma for convenience. Let K and $\tilde{K}$ be compact Jordan domains in general position. Let $z_1, z_2, z_3 \in \partial K \setminus \partial \tilde{K}$ appear in counterclockwise order, similarly $\tilde{z}_1, \tilde{z}_2, \tilde{z}_3 \in \partial \tilde{K} \setminus \partial K$. An indexable homeomorphism $f : \partial K \to \partial \tilde{K}$ is called *faithful* if it sends $z_i \mapsto \tilde{z}_i$ for $i = 1, 2, 3$. We wish to find a faithful f with $\eta(f) \geq 0$. We refer to the z_i and the $\tilde{z}_i$ as our *constraint points*.

We proceed by induction on the number of intersection points $\partial K \cap \partial \tilde{K}$, noting that this number is always even. The Circle Index Lemma takes care of the cases where ∂K and $\partial \tilde{K}$ meet 0 or 2 times. Thus suppose that ∂K and $\partial \tilde{K}$ meet at least 4 times.

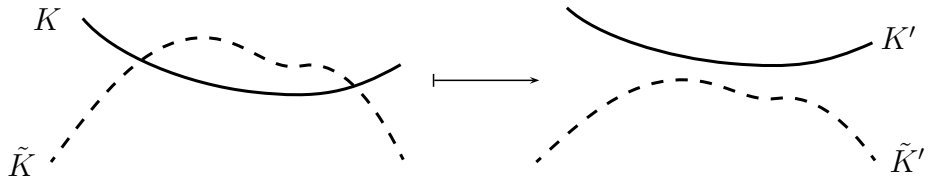

Figure 6.13: **"Pulling apart" two sub-curves of ∂K and $\partial \tilde{K}$ "crossing minimally."** Note that *a priori* the orientations on ∂K and $\partial \tilde{K}$ may be arbitrary, so the open Jordan domain formed between the two curves on the left may be contained in both K and $\tilde{K}$, in just one of them, or in neither.

The main idea of the proof is to find sub-arcs of ∂K and $\partial \tilde{K}$ which "cross minimally," and "pull them apart," see Figure 6.13. When "pulling these sub-arcs apart," we leave the rest of ∂K and $\partial \tilde{K}$ fixed, as in Figure 6.14. We then apply the induction hypothesis to the resulting Jordan domains K' and $\tilde{K}'$, obtaining some indexable homeomorphism $f' : \partial K' \to \partial \tilde{K}'$ with $\eta(f') \geq 0$. Finally, we use f' to construct f, arguing that the fixed-point index is preserved or increased in this last construction. The details of the proof will be worked out in a torus parametrization. This is because the torus parametrization allows us to work systematically and somewhat combinatorially through many cases.

6.5.1 Initial set-up

More precisely, two intersection points $p_0, \tilde{p}_0 \in \partial K \cap \partial \tilde{K}$ are called *adjacent in* ∂K if one of the two arcs $[p_0 \to \tilde{p}_0]_{\partial K}$ and $[\tilde{p}_0 \to p_0]_{\partial K}$ does not contain any other intersection points $\partial K \cap \partial \tilde{K}$. Note that two intersection points of $\partial K \cap \partial \tilde{K}$ may be adjacent only if ∂K enters $\tilde{K}$ at one and leaves $\tilde{K}$ at the other, so our notation is in keeping with convention. In particular we will suppose that ∂K enters $\tilde{K}$ at p_0 and exits $\tilde{K}$ at $\tilde{p}_0$. We define *adjacency in* $\partial \tilde{K}$ similarly.

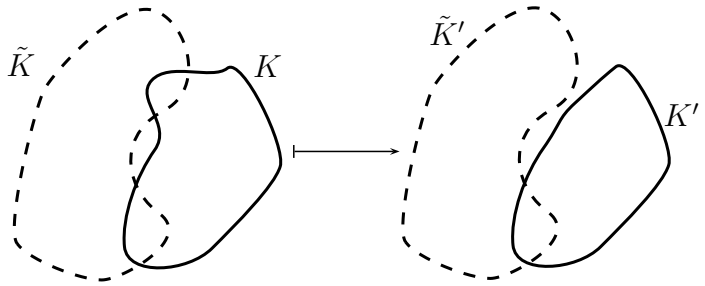

Figure 6.14: **The compact Jordan domains under consideration, before and after "pulling the sub-arcs apart."** Note that in this case, we had three choices of which arcs to pull apart, and each choice results in a different pair $K', \tilde{K}'$.

Two intersection points $p_0, \tilde{p}_0 \in \partial K \cap \partial \tilde{K}$ are called *doubly adjacent* if they are adjacent both in ∂K and in $\partial \tilde{K}$. Then the notion of sub-arcs of ∂K and $\partial \tilde{K}$ "crossing minimally" which we informally described before is exactly captured by the property of double adjacency. The following is easy to see by induction on the number of intersection points $\partial K \cap \partial \tilde{K}$:

Observation 6.17. *If $\partial K \cap \partial \tilde{K}$ meet, then there are doubly adjacent pairs of points in $\partial K \cap \partial \tilde{K}$.*

From now on, we fix a torus parametrization of ∂K and $\partial \tilde{K}$ in $\mathbb{T} = \mathbb{S}^1 \times \mathbb{S}^1$ via $\kappa : \partial K \to \mathbb{S}^1$ and $\tilde{\kappa} : \partial \tilde{K} \to \mathbb{S}^1$. Let $p_0, \tilde{p}_0$ be doubly adjacent. Then we denote $P_0 = (\kappa(p_0), \tilde{\kappa}(p_0))$ and $\tilde{P}_0 = (\kappa(\tilde{p}_0), \tilde{\kappa}(\tilde{p}_0))$ as usual, so $P_0 \in \mathbb{T}$ parametrizes the point $p_0 \in \partial K \cap \partial \tilde{K}$, similarly $\tilde{P}_0$ for $\tilde{p}_0$. We can now see how we plan to apply induction:

Observation 6.18. *Let the torus parametrization $\mathbb{T}'$ be obtained from $\mathbb{T}$ by deleting P_0 and $\tilde{P}_0$. Then $\mathbb{T}'$ parametrizes a pair $\partial K'$ and $\partial \tilde{K}'$, where K' and $\tilde{K}'$ are compact Jordan domains in general position, meeting two times fewer than do K and $\tilde{K}$.*

Thus our strategy will be as follows: first, apply the induction hypothesis to $\mathbb{T}'$, to get a faithful indexable γ' so that $w(\gamma') = \eta(f') \geq 0$; then, "reinsert" the points P_0 and $\tilde{P}_0$, in such a way that the fixed-point index is preserved or increased.

6.5.2 Partitioning $\mathbb{T}$

We will refer to the collection of the sets $\mathbb{S}^1 \times \{\tilde{\kappa}(\tilde{z}_i)\}$ and $\{\kappa(z_i)\} \times \mathbb{S}^1$ for $i = 1, 2, 3$ as our *grid lines*, and to their pairwise intersection points as our *lattice points*. Furthermore we refer to the points $(\kappa(z_i), \tilde{\kappa}(\tilde{z}_i))$, for $i = 1, 2, 3$, as our *constraint lattice points*. If we draw our torus parametrization with a lattice point chosen to be

the base point, then the grid lines actually divide the drawing into a grid, with nine disjoint rectangular cells in total. We will refer to these as our *lattice cells*. Within each lattice cell there are well-defined left, right, up, and down directions, in the same way as for the adjacency box A.

From now on, whenever we draw the torus parametrization, we will always pick one of the lattice points for the base point. Suppose for the rest of this paragraph that we have fixed such a drawing. Then we may refer to the lattice cells according to their position in this drawing in the natural way, for example we may refer to the bottom-left lattice cell. We denote the lattice cells by $C_\uparrow, C_\nearrow, C_\rightarrow, \ldots, C_\leftarrow, C_\nwarrow$ in the natural way, and write $C_\bullet$ to denote the central lattice cell. Then the union of the closed upper-left, middle-left, and lower-left lattice cells is called the *left lattice column*. We define the *central* and *right lattice columns* analogously, as well as the *bottom*, *central*, and *top lattice rows*. We again emphasize that these descriptors depend crucially on the choice of base point for the drawing of the torus parametrization.

As usual, we denote $P_i = (\kappa(p_i), \tilde{\kappa}(p_i))$ and $\tilde{P}_i = (\kappa(\tilde{p}_i), \tilde{\kappa}(\tilde{p}_i))$, where $\{p_i, \tilde{p}_i\} = \partial K \cap \partial \tilde{K}$, with the same notational conventions as usual. We will apply Lemma 6.1 frequently. Before moving on to the next part of the proof, we make an observation that will simplify things later. This also serves as a good warm-up to and refresher on Lemma 6.1 and the torus parametrization.

Observation 6.19. *Suppose that all of the P_i and $\tilde{P}_i$ lie in a single lattice row, or in a single lattice column. Then Lemma 3.5 holds.*

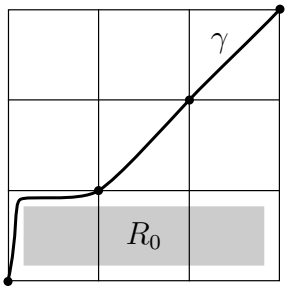

Figure 6.15: **The situation if all P_i and $\tilde{P}_i$ lie in a single lattice row.** Here we have supposed that they lie in the bottom lattice row, in which case they in fact lie in some open rectangle R_0 contained in the bottom lattice row.

To see why, we walk through Figure 6.15. We suppose that all of the P_i and $\tilde{P}_i$ lie in a single lattice row, call it R. The same argument will work if they all lie in a single lattice column. Without loss of generality, we have drawn the torus parametrization so that R is the bottom lattice row, and so that the base point

is a constraint lattice point. Then the other two constraint lattice points are the ones along the diagonal, drawn in the figure as dots. As usual, a faithful indexable homeomorphism $f : \partial K \to \partial \tilde{K}$ can be parametrized by a "strictly increasing" curve

- from the bottom-left corner of the bottom-left cell,
- to the top-right corner of the top-right cell,
- passing through the two other constraint lattice points, and
- missing all of the P_i and $\tilde{P}_i$.

Thus the γ we have drawn in Figure 6.15 parametrizes such a homeomorphism. Furthermore, because there are finitely many P_i and $\tilde{P}_i$, and all of them lie in the bottom lattice row R, in fact they lie in an open rectangle R_0 as depicted in Figure 6.15. Thus $\eta(f) \geq 0$ by Lemma 6.1.

6.5.3 The adjacency box

Let p_0 and $\tilde{p}_0$ be adjacent in ∂K, and let a_0 denote the arc connecting p_0 and $\tilde{p}_0$ along ∂K which contains no other intersection points $\partial K \cap \partial \tilde{K}$. Thus a_0 is equal to one of $[p_0 \to \tilde{p}_0]_{\partial K}$ and $[\tilde{p}_0 \to p_0]_{\partial K}$. There is a unique such a_0 so long as ∂K and $\partial \tilde{K}$ meet at least four times, as is our running assumption. We call this a_0 the *short adjacency arc of* $p_0, \tilde{p}_0$ *in* ∂K, and define the *short adjacency arc* $\tilde{a}_0$ *of* $p_0, \tilde{p}_0$ *in* $\partial \tilde{K}$ similarly.

Continuing with the notation of the last paragraph, let a be obtained from a_0 by extending it slightly on both ends, so that $a \subset \partial K$ is a closed Jordan arc containing p_0 and $\tilde{p}_0$, but no other points of $\partial K \cap \partial \tilde{K}$, and no constraint point z_i which was not already contained in a_0. There is a topologically unique such a. Then we say that a is the *adjacency arc of* $p_0, \tilde{p}_0$ *in* ∂K of $p_0, \tilde{p}_0$. We define the *adjacency arc* $\tilde{a}$ *of* $p_0, \tilde{p}_0$ *in* $\partial \tilde{K}$ similarly.

Again continuing with the notation of before, let $A = \kappa(a) \times \tilde{\kappa}(\tilde{a}) \subset \mathbb{T} = \mathbb{S}^1 \times \mathbb{S}^1$ be the closed rectangle so that $(x, \tilde{x}) \in A$ if and only if $\kappa^{-1}(x) \in a$ and $\tilde{\kappa}^{-1}(\tilde{x}) \in \tilde{a}$. There is a topologically unique such A. Then A is called the *adjacency box* of $p_0, \tilde{p}_0$.

6.5.4 Enumerating the possible locations of the adjacency box

Pick doubly adjacent $p_0, \tilde{p}_0$, having adjacency arcs $a_0, \tilde{a}_0$ and let $A = \kappa(a_0) \times \tilde{\kappa}(\tilde{a}_0) \subset \mathbb{T}$ be their adjacency box. Let $b_0 = \partial K \setminus a_0$ and $\tilde{b}_0 = \partial \tilde{K} \setminus b_0$, and set $B = \kappa(b_0) \times \tilde{\kappa}(\tilde{b}_0) \subset \mathbb{T}$. Then b_0 and $\tilde{b}_0$ are called the *blotchy arcs* of $p_0, \tilde{p}_0$ in ∂K and $\partial \tilde{K}$, respectively, and B is called the *blotchy box* of $p_0, \tilde{p}_0$. This is because of the following observation:

Observation 6.20. *The parametrization in $\mathbb{T}$ of every point in $\partial K \cap \partial \tilde{K} \setminus \{p_0, \tilde{p}_0\}$ lies in the blotchy box B of $p_0, \tilde{p}_0$. However, we have basically no* a priori *information about how these points are arranged in B.*

Fix doubly adjacent $p_0, \tilde{p}_0$, having adjacency arcs $a_0, \tilde{a}_0$, adjacency box A, blotchy arcs $b_0, \tilde{b}_0$, and blotchy box B.

Observation 6.21. *Suppose that $z_1, z_2, z_3 \in a_0$. Then $b_0 \times \mathbb{S}^1 \supset B$ is contained in a lattice column, regardless of which lattice point is chosen for the base point of our drawing of the torus parametrization. Thus Lemma 3.5 holds by Observation 6.19.*

Thus we may assume without loss of generality that $z_1 \notin a_0$. Similarly, one of the $\tilde{z}_i$ lies outside of $\tilde{a}_0$, but we cannot say which after having fixed the labeling of the z_i and $\tilde{z}_i$ by insisting that $z_1 \notin a_0$. Next:

Observation 6.22. *We may suppose without loss of generality that A meets the left lattice column.*

To see why, recall that we have insisted that the z_1, z_2, z_3 appear in that order as we traverse ∂K positively. Thus traverse ∂K starting from z_1 positively, until arriving at an endpoint of a_0. Let z_i be the last one of z_1, z_2, z_3 crossed during this traversal. Then we may draw the torus parametrization using $(\kappa(z_i), \tilde{\kappa}(\tilde{z}_i))$ as the base point.

After relabeling, we may suppose that our torus parametrization is drawn with $(\kappa(z_1), \tilde{\kappa}(\tilde{z}_1))$ as our base point. Then:

Observation 6.23. *Under our running assumptions, the topological location of A in the torus parametrization is completely determined by two pieces of information:*

- *which of the cells of the left lattice column contains the lower-left corner $A_{\swarrow}$ of A, and*
- *which of the nine lattice cells contains the upper-right corner $A_{\nearrow}$ of A.*

We will denote the cases for the topological location of A in the following way: let r equal one of $\swarrow, \leftarrow, \nwarrow$, and let s equal one of $\uparrow, \nearrow, \rightarrow, \ldots, \leftarrow, \nwarrow, \bullet$. Then we say that case (r, s) occurs if and only if $A_{\swarrow}$ lies in lattice cell C_r, and $A_{\nearrow}$ lies in lattice cell C_s. *A priori* this gives us $3 \times 9 = 27$ possibilities for A. Fortunately we will see that many of the cases are handled in more or less the same way. For reference we have depicted the 27 cases in Figure 6.16.

For future reference, we isolate and note an observation limiting the possibilities for A which we need to consider:

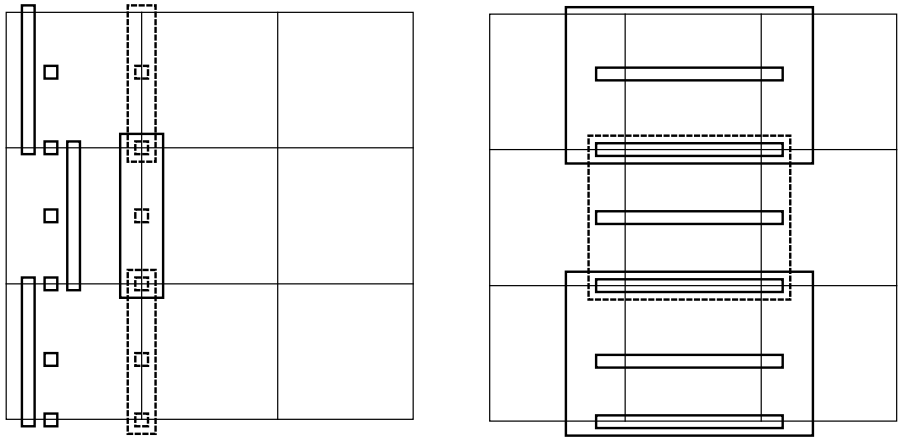

(a) (b)

Figure 6.16: **The 27 possible topologically distinct locations for the adjacency box A.** Both torus parametrizations are drawn with $(\kappa(z_1), \tilde{\kappa}(\tilde{z}_1))$ as the base point. In (a) we have cases (r, s) where $t \neq \nearrow, \rightarrow, \searrow$. The adjacency box A is shown with a solid boundary for the cases $s = \nwarrow, \leftarrow, \swarrow$, and with a dashed boundary for the cases $r = \uparrow, \bullet, \downarrow$, except for the case $(\swarrow, \uparrow)$, which is shown with a solid boundary. In (b), we have cases (r, s) where $s = \nearrow, \rightarrow, \searrow$. Here the adjacency box A is drawn with a solid boundary, except for the case $(\swarrow, \nearrow)$, which is shown with a dashed boundary.

Observation 6.24. *We suppose without loss of generality that the adjacency box A does not meet the left/right grid line $\{\kappa(z_1)\} \times \mathbb{S}^1$.*

6.5.5 First application of induction, when γ misses A

For the remainder of the proof, let γ be obtained inductively as described in Observation 6.18. We wish to compare $w(\gamma)$ with $w(\gamma')$. The germane issue is the effect of the presence of $P_0, \tilde{P}_0$ in $\mathbb{T}$ on $w(\gamma)$ compared to their absence from $\mathbb{T}'$ when computing $w(\gamma')$. In light of Lemma 6.1, there are precisely two ways that reinserting these points can change the fixed-point index of $\gamma = \gamma'$:

- First, they may affect the winding numbers $\omega(\partial K, \tilde{z}_1)$ and $\omega(\partial \tilde{K}, z_1)$ as compared to $\omega(\partial K, \tilde{z}_1)$ and $\omega(\partial \tilde{K}, z_1)$, and
- second, they may affect the sum $-\#p_\downarrow(z_1, \gamma) + \#\tilde{p}_\downarrow(z_1, \gamma)$, equivalently the sum $\#p_\uparrow(z_1, \gamma) - \#\tilde{p}_\uparrow(z_1, \gamma)$, as compared to the respective sums in $\mathbb{T}'$.

We have the following nice observation:

Observation 6.25. *The winding numbers $\omega(\partial K, \tilde{z}_1)$ and $\omega(\partial K', \tilde{z}_1')$ are equal.*

This follows from Observation 6.24. In particular, it follows because $\omega(\partial K, \tilde{z}_1)$ is determined completely by whether we arrive first at a $\kappa(p_i)$ or a $\kappa(\tilde{p}_i)$ as we traverse the first coordinate $\mathbb{S}^1$ of $\mathbb{T}$ positively starting at $\kappa(z_1)$, and this is invariant under reinserting $P_0, \tilde{P}_0$ by Observation 6.24.

Next:

Claim 6.26. *Suppose that $\gamma = \gamma'$ does not meet A. Then $w(\gamma) = w(\gamma') \geq 0$.*

Proof. There are two cases, handled differently:

Case 1. The box A does not meet the bottom/top grid line $\mathbb{S}^1 \times \{\tilde{\kappa}(\tilde{z}_1)\}$.

In this case, similarly to Observation 6.25, we get that $\omega(\partial \tilde{K}, z_1) = \omega(\partial \tilde{K}', z_1')$. Also, because $\gamma = \gamma'$ does not meet A, we have that P_0 and $\tilde{P}_0$ either both lie above γ in $\mathbb{T}$, or both lie below it. Thus $-\#p_\downarrow(z_1, \gamma) + \#\tilde{p}_\downarrow(z_1, \gamma) = -\#p_\downarrow(z_1, \gamma') + \#\tilde{p}_\downarrow(z_1, \gamma')$, completing the proof of the claim in this case.

Case 2. The box A crosses the bottom/top grid line $\mathbb{S}^1 \times \{\tilde{\kappa}(\tilde{z}_1)\}$.

In this case, the bottom/top grid line $\mathbb{S}^1 \times \{\tilde{\kappa}(\tilde{z}_1)\}$ cuts A into two connected components $A_\downarrow$ and $A_\uparrow$, one of which contains P_0 and the other of which contains $\tilde{P}_0$. The "upper half" of A, which we have denoted $A_\uparrow$, actually lies in the bottom lattice row of $\mathbb{T}$, and furthermore lies beneath γ because γ does not meet A. Similarly $A_\downarrow$ lies in the upper lattice row of $\mathbb{T}$ and above γ. There are two sub-cases:

Sub-case 2.1. $P_0 \in A_\downarrow$, $\tilde{P}_0 \in A_\uparrow$

In this case $\omega(\partial \tilde{K}, z_1) = 1$ and $\omega(\partial \tilde{K}', z_1') = 0$. On the other hand $-\#p_\downarrow(z_1, \gamma) + \#\tilde{p}_\downarrow(z_1, \gamma)$ is 1 less than $-\#p_\downarrow(z_1, \gamma') + \#\tilde{p}_\downarrow(z_1, \gamma')$. Thus $w(\gamma) = w(\gamma')$ by Lemma 6.1.

Sub-case 2.2. $P_0 \in A_\uparrow$, $\tilde{P}_0 \in A_\downarrow$

This sub-case is handled in the same way as Sub-case 2.1, except this time $\omega(\partial \tilde{K}, z_1) + 1 = \omega(\partial \tilde{K}', z_1')$ and $-\#p_\downarrow(z_1, \gamma) = -\#p_\downarrow(z_1, \gamma') + 1$. This completes the proof of Claim 6.26. □

Much of the rest of the proof of Lemma 3.5 will be spent trying to reduce to the situation handled by Claim 6.26.

6.5.6 Moving $P_0, \tilde{P}_0$, or adjusting A

We continue with the notation from before. Because A is a topological rectangle in the natural way, it has well-defined left, right, up, and down directions. More precisely, we suppose that the positive orientation on the first coordinate of $\mathbb{T} = \mathbb{S}^1 \times \mathbb{S}^1$ goes left-to-right, and on the second coordinate goes down-to-up. *A priori* there are exactly four topologically distinct ways that P_0 and $\tilde{P}_0$ may be arranged in the adjacency box A of $p_0, \tilde{p}_0$, as depicted in Figure 6.17.

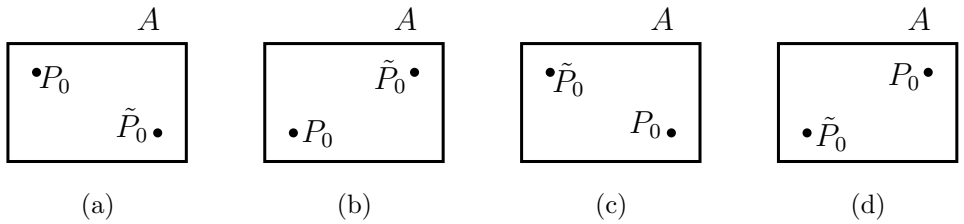

Figure 6.17: **The *a priori* possible arrangements of $P_0, \tilde{P}_0$ in the adjacency box B.**

The following observation will be helpful later:

Observation 6.27. *We may move P_0 and $\tilde{P}_0$ as we please within A, without changing the topological configuration of K, $\tilde{K}$, and the constraint points, so long as the following two requirements are satisfied:*

- *Each of P_0 and $\tilde{P}_0$ must remain in the same respective connected component of $A \setminus \cup\{\text{grid lines}\}$.*
- *We must not change which of the cases of Figure 6.17 occurs.*

During the eventual induction step, it will often be desirable to move the points P_0 and $\tilde{P}_0$. By Observation 6.27 we may do so, within limits, without affecting any germane aspects of the situation. It will also be useful to note the following:

Observation 6.28. *The points P_0 and $\tilde{P}_0$ lie in different connected components of $A \setminus \cup\{\text{grid lines}\}$, unless $A = A \setminus \cup\{\text{grid lines}\}$.*

This is by our construction of A.

Next, recall that we had latitude in choosing the adjacency box A of $p_0, \tilde{p}_0$. With this in mind, we make the following observation:

Observation 6.29. *We may replace A by any topological rectangle contained in A, having sides "parallel to" those of A in the natural sense, so long as the interior of*

this new rectangle continues to contain P_0 and $\tilde{P}_0$. This replacement preserves the germane features of A in that every argument we make continues to work under this replacement.

We now apply these observations to describe how to complete the proof in a crucial case:

Claim 6.30. *Suppose that γ meets only a single connected component of the set $A \setminus \cup\{\text{grid lines}\}$, and furthermore that this component contains a corner of A. Then we may reduce to Claim 6.26, completing the proof of Lemma 3.5.*

Proof. First note that $\gamma \cap A$ has only a single connected component, because γ is "strictly increasing" in $\mathbb{T}$. This also implies that if the component of $A \setminus \cup\{\text{grid lines}\}$ which meets γ contains a corner of A, then it must contain either the upper-left corner, or the lower-right corner, or both. Suppose it contains the upper-left corner. The argument will be the same if it contains the lower-right corner. A potential drawing of the situation is depicted in Figure 6.18.

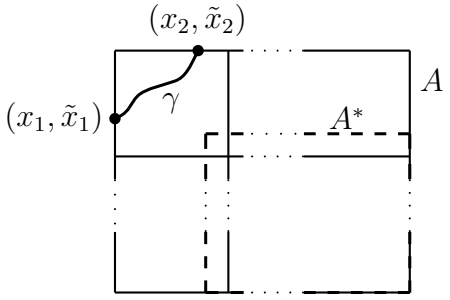

Figure 6.18: **Moving P_0 and $\tilde{P}_0$ to find a replacement for A not meeting γ.**

Then it is clear that within the constraints described in Observation 6.27 it is possible to move P_0 and $\tilde{P}_0$ if necessary so that both lie to the right of $(x_2, \tilde{x}_2)$ and below $(x_1, \tilde{x}_1)$, and then it is clear how to choose A^* our replacement rectangle for A. This completes the proof of Claim 6.30. □

6.5.7 Restricting the possibilities further

Let $p_0, \tilde{p}_0$ be doubly adjacent, with short adjacency arcs a_0 and $\tilde{a}_0$. The unoriented union $a_0 \cup \tilde{a}_0$ is an unoriented Jordan curve. The orientations on a_0 and $\tilde{a}_0$ may or may not agree, and may or may not induce the positive orientation on $a_0 \cup \tilde{a}_0$. Let U be the open Jordan domain bounded by $a_0 \cup \tilde{a}_0$. Then U is called the *adjacency domain* of $p_0, \tilde{p}_0$. We make the following observation:

Observation 6.31. *There exist at least two distinct pairs of doubly adjacent points in $\partial K \cap \partial \tilde{K}$ whose adjacency domains are contained in $\tilde{K}$.*

This is easy to show by induction, similarly to Observation 6.17. Then we get the following as a corollary:

Observation 6.32. *There exist at least two distinct pairs of doubly adjacent points in $\partial K \cap \partial \tilde{K}$, so that if $p_0, \tilde{p}_0$ is such a pair, then their short adjacency arc in ∂K is $a_0 = [p_0 \to \tilde{p}_0]_{\partial K}$.*

We may restate Observation 6.32 in the following way:

Observation 6.33. *There exist at least two distinct pairs of doubly adjacent points in $\partial K \cap \partial \tilde{K}$, so that if $P_0, \tilde{P}_0 \in \mathbb{T}$ parametrize one of these pairs, then P_0 lies to the left of $\tilde{P}_0$ in their adjacency box A.*

Thus without loss of generality, for the remainder of the proof, we restrict our attention to the arrangements in Figures 6.17a, 6.17b.

6.5.8 Completing the proof if a constraint lattice point lies in the adjacency box A

Claim 6.34. *Suppose that A contains a constraint lattice point. Then we may move P_0 and $\tilde{P}_0$ around in such a way that we get $w(\gamma) \geq 0$, completing the proof of Lemma 3.5.*

Proof. The following will be useful to keep in mind:

Observation 6.35. *The curve γ cannot cross grid lines, except at constraint lattice points.*

We prove Claim 6.34 in four cases:

Case 1. The arrangement in Figure 6.17a occurs, and A does not meet the bottom/top grid line of $\mathbb{T}$.

It is not hard to see that then P_0 and $\tilde{P}_0$ are on different sides of γ in our drawing of $\mathbb{T}$, and furthermore that P_0 lies above γ and $\tilde{P}_0$ below it. Then, arguing straightforwardly as in the proof of Claim 6.26, Case 1, we get $w(\gamma) = w(\gamma') + 1$.

For the remaining cases, the following will be the key observation:

Observation 6.36. *Suppose that P_0 lies in a connected component of $A \backslash \cup\{\text{grid lines}\}$ through which γ passes. Then we may decide whether P_0 lies above or below γ in our drawing of $\mathbb{T}$, in the sense that we may move P_0 within the constraints given in Observation 6.27 to arrange either situation. The same is true for $\tilde{P}_0$.*

Case 2. The arrangement in Figure 6.17a occurs, and A crosses the bottom/top grid line of $\mathbb{T}$.

We argue similarly to the proof of Claim 6.26, Case 2. In this case it can always be arranged that P_0 lies below γ and that $\tilde{P}_0$ lies above it. There are three cases to check, and in each it is clear how to proceed. The cases are shown in Figure 6.19. Again we get $w(\gamma) = w(\gamma') + 1$.

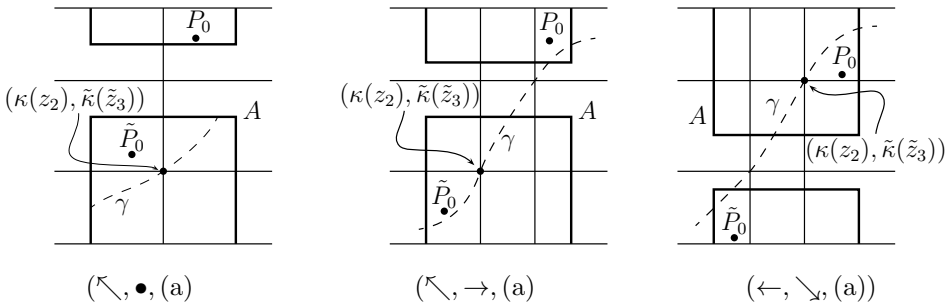

Figure 6.19: **The three possibilities for A when A contains a constraint lattice point and meets the bottom/top grid line, and Figure 6.17a occurs.**

Case 3. The arrangement in Figure 6.17b occurs, and A does not meet the bottom/top grid line of $\mathbb{T}$.

As per our usual arguments, we have $\omega(\partial\tilde{K}, z_1) = \omega(\partial\tilde{K}', z_1')$ and $\omega(\partial K, \tilde{z}_1) = \omega(\partial K', \tilde{z}_1')$. Thus we will be done if we can argue that we may arrange so that P_0 and $\tilde{P}_0$ lie on the same side of γ in our drawing of $\mathbb{T}$. There are five cases to check, and in each it is clear how to proceed. Figure 6.20 shows each case.

Case 4. The arrangement in Figure 6.17b occurs, and A crosses the bottom/top grid line of $\mathbb{T}$.

In this case we have that $\omega(\partial K, \tilde{z}_1) = \omega(\partial K', \tilde{z}_1')$ as per Observation 6.25, but unfortunately $\omega(\partial\tilde{K}, z_1) = \omega(\partial\tilde{K}', z_1') - 1$. Thus we will be done if we can argue that we may arrange so that P_0 lies above γ, and $\tilde{P}_0$ below it, in our drawing of $\mathbb{T}$. There are only three cases to check. Figure 6.21 shows each. This completes the proof of Claim 6.34. □

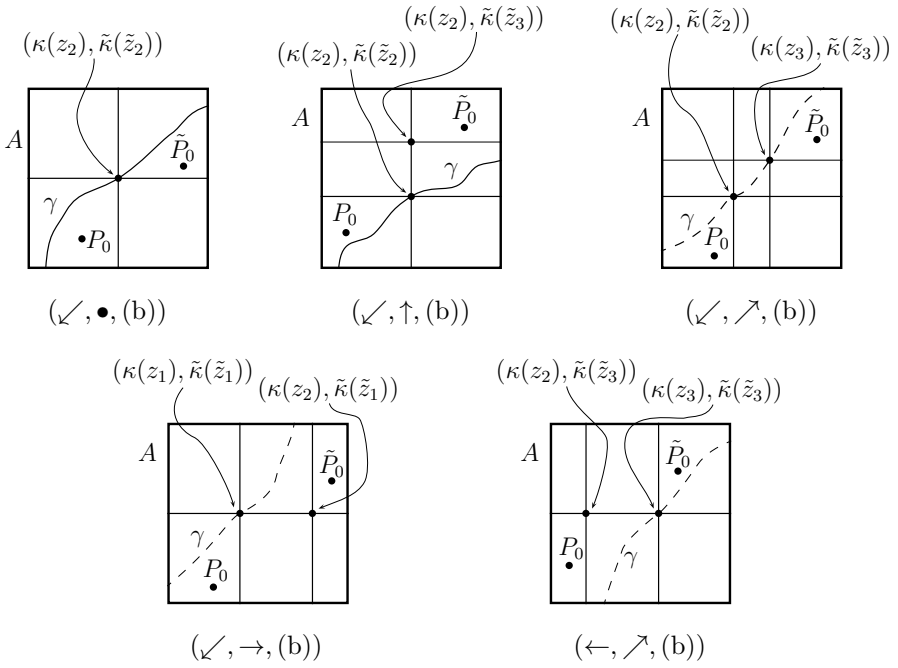

Figure 6.20: **The five possibilities for A when A contains a constraint lattice point and does not meet the bottom/top grid line, and Figure 6.17b occurs.**

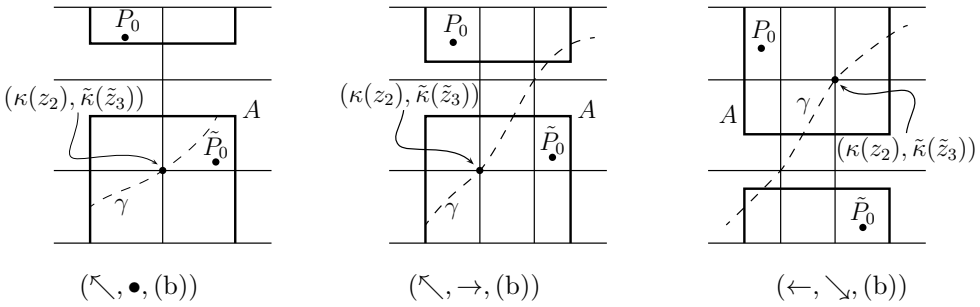

Figure 6.21: **The three possibilities for A when A contains a constraint lattice point and does not meet the bottom/top grid line, and Figure 6.17b occurs.**

6.5.9 Completing the proof of Lemma 3.5

We summarize the situation so far. There are $3 \times 9 = 27$ possibilities for the location of A in $\mathbb{T}$. Most of the cases fall into one of the following three categories. For a given type (s, t) of A, it is quick and easy to check the conditions that define membership in the categories. We also refer the reader back to Figure 6.16, where each of the 27 cases is drawn.

Category 1. The intersection $A \cap (C_{\swarrow} \cup C_{\bullet} \cup C_{\nearrow})$ is empty.

Then we are done by Claim 6.26, because γ is restricted to lie in $C_{\swarrow} \cup C_{\bullet} \cup C_{\nearrow}$. Next:

Category 2. The set $A \setminus \cup\{\text{grid lines}\}$ has exactly one connected component lying in $C_{\swarrow} \cup C_{\bullet} \cup C_{\nearrow}$, and this connected component contains a corner of A.

Then we are done by Claim 6.30, again because γ is restricted to lie in $C_{\swarrow} \cup C_{\bullet} \cup C_{\nearrow}$. Finally:

Category 3. The adjacency box A contains a constraint lattice point.

This is exactly handled by Claim 6.34.

In Figure 6.22 we list the 27 possible locations of A, and for each falling into one of these three categories, we indicate which.

r	s	Category
$\nwarrow$	$\uparrow$	1
$\nwarrow$	$\nearrow$	3
$\nwarrow$	$\rightarrow$	2
$\nwarrow$	$\searrow$	3
$\nwarrow$	$\downarrow$	3
$\nwarrow$	$\swarrow$	3
$\nwarrow$	$\leftarrow$	
$\nwarrow$	$\nwarrow$	1
$\nwarrow$	$\bullet$	2

r	s	Category
$\leftarrow$	$\uparrow$	3
$\leftarrow$	$\nearrow$	2
$\leftarrow$	$\rightarrow$	
$\leftarrow$	$\searrow$	2
$\leftarrow$	$\downarrow$	
$\leftarrow$	$\swarrow$	3
$\leftarrow$	$\leftarrow$	1
$\leftarrow$	$\nwarrow$	1
$\leftarrow$	$\bullet$	3

r	s	Category
$\swarrow$	$\uparrow$	2
$\swarrow$	$\nearrow$	2
$\swarrow$	$\rightarrow$	2
$\swarrow$	$\searrow$	3
$\swarrow$	$\downarrow$	3
$\swarrow$	$\swarrow$	3
$\swarrow$	$\leftarrow$	3
$\swarrow$	$\nwarrow$	3
$\swarrow$	$\bullet$	2

Figure 6.22: **The cases to check for Lemma 3.5.**

Case $(\leftarrow, \downarrow)$ is easily handled in a manner similar to the proof of Claim 6.30. If case $(\leftarrow, \rightarrow)$ occurs, then our running assumption that P_0 lies to the left of $\tilde{P}_0$ in A implies that $w(\gamma) = w(\gamma') + 1$ by Lemma 6.1 as usual.

We have no way to deal with case $(\nwarrow, \leftarrow)$ directly. However, recall from Observation 6.33 that we had two choices for $P_0, \tilde{P}_0$ for which every step of the proof until this point go through. Because the adjacency arc $\tilde{a}_0$ contains two constraint points $\tilde{z}_i$ if case $(\nwarrow, \leftarrow)$ occurs, it can occur for only one of these two choices of $P_0, \tilde{P}_0$, so we do not need to handle this case directly. This completes the proof of Lemma 3.5. $\square$

Appendix: The relevant topological configurations of three-disk subsets of $\{A, B, \tilde{A}, \tilde{B}\}$

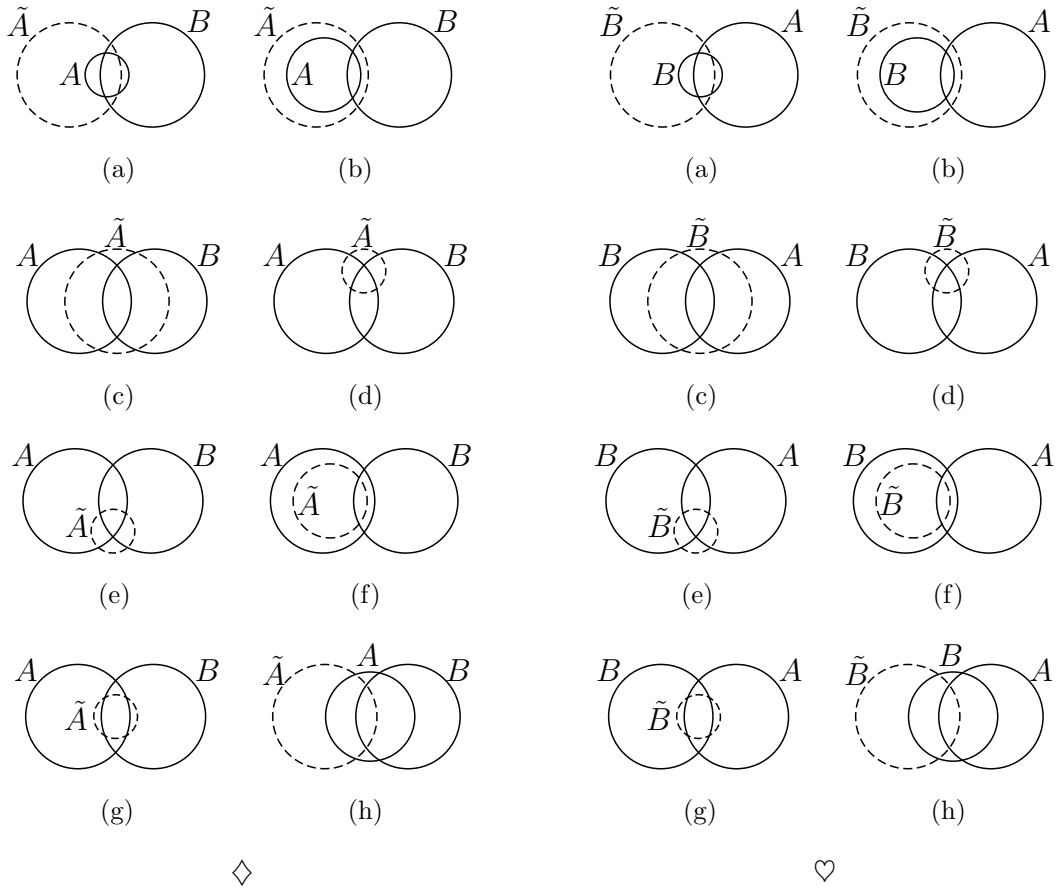

Figures ♢, ♡: **The possible topological configurations of $\{A, B, \tilde{A}\}$ and $\{A, B, \tilde{B}\}$ respectively, under the hypotheses of Proposition 5.11.**

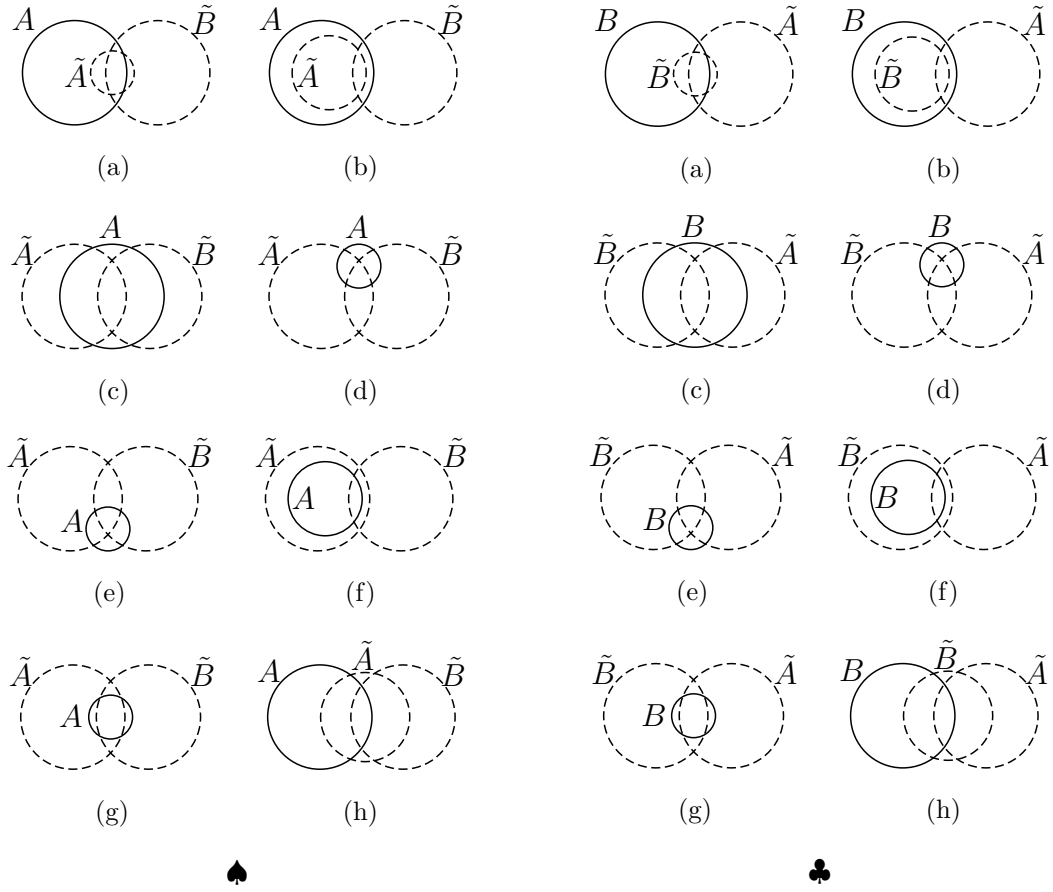

♠ ♣

Figures ♠, ♣: **The possible topological configurations of $\{A, \tilde{A}, \tilde{B}\}$ and $\{B, \tilde{A}, \tilde{B}\}$ respectively, under the hypotheses of Proposition 5.11.**

References

[Ahl78] Lars V. Ahlfors, *Complex analysis*, 3rd ed., McGraw-Hill Book Co., New York, 1978. An introduction to the theory of analytic functions of one complex variable; International Series in Pure and Applied Mathematics. MR510197 (80c:30001)

[And70] E. M. Andreev, *Convex polyhedra of finite volume in Lobačevskiĭ space*, Mat. Sb. (N.S.) **83 (125)** (1970), 256–260 (Russian). MR0273510 (42 #8388)

[BS90] Alan F. Beardon and Kenneth Stephenson, *The uniformization theorem for circle packings*, Indiana Univ. Math. J. **39** (1990), no. 4, 1383–1425, DOI 10.1512/iumj.1990.39.39062. MR1087197 (92b:52038)

[BS91a] Alan F. Beardon and Kenneth Stephenson, *Circle packings in different geometries*, Tohoku Math. J. (2) **43** (1991), no. 1, 27–36, DOI 10.2748/tmj/1178227533. MR1088712 (92a:52023)

[BS91b] Alan F. Beardon and Kenneth Stephenson, *The Schwarz-Pick lemma for circle packings*, Illinois J. Math. **35** (1991), no. 4, 577–606. MR1115988 (93a:30028a)

[BS04] Alexander I. Bobenko and Boris A. Springborn, *Variational principles for circle patterns and Koebe's theorem*, Trans. Amer. Math. Soc. **356** (2004), no. 2, 659–689, DOI 10.1090/S0002-9947-03-03239-2. MR2022715 (2005b:52054)

[BP92] Riccardo Benedetti and Carlo Petronio, *Lectures on hyperbolic geometry*, Universitext, Springer-Verlag, Berlin, 1992. MR1219310 (94e:57015)

[Brä92] Walter Brägger, *Kreispackungen und Triangulierungen*, Enseign. Math. (2) **38** (1992), no. 3-4, 201–217 (German). MR1189006 (94b:52032)

[BS93] Graham R. Brightwell and Edward R. Scheinerman, *Representations of planar graphs*, SIAM J. Discrete Math. **6** (1993), no. 2, 214–229, DOI 10.1137/0406017. MR1215229 (95d:05043)

[Bro86] Robert Brooks, *Circle packings and co-compact extensions of Kleinian groups*, Invent. Math. **86** (1986), no. 3, 461–469, DOI 10.1007/BF01389263. MR860677 (88b:32050)

[Car13] C. Carathéodory, *Über die gegenseitige Beziehung der Ränder bei der konformen Abbildung des Inneren einer Jordanschen Kurve auf einen Kreis*, Math. Ann. **73** (1913), no. 2, 305–320, DOI 10.1007/BF01456720 (German). MR1511735

[CdV91] Yves Colin de Verdière, *Un principe variationnel pour les empilements de cercles*, Invent. Math. **104** (1991), no. 3, 655–669, DOI 10.1007/BF01245096 (French). MR1106755 (92h:57020)

[FM12] Benson Farb and Dan Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012. MR2850125

[For91] Otto Forster, *Lectures on Riemann surfaces*, Graduate Texts in Mathematics, vol. 81, Springer-Verlag, New York, 1991. Translated from the 1977 German original by Bruce Gilligan; Reprint of the 1981 English translation. MR1185074 (93h:30061)

[He90] Zheng-Xu He, *Solving Beltrami equations by circle packing*, Trans. Amer. Math. Soc. **322** (1990), no. 2, 657–670, DOI 10.2307/2001719. MR974518 (91c:30032)

[He99] Zheng-Xu He, *Rigidity of infinite disk patterns*, Ann. of Math. (2) **149** (1999), no. 1, 1–33, DOI 10.2307/121018. MR1680531 (2000j:30068)

[HS93] Zheng-Xu He and Oded Schramm, *Fixed points, Koebe uniformization and circle packings*, Ann. of Math. (2) **137** (1993), no. 2, 369–406, DOI 10.2307/2946541. MR1207210 (96b:30015)

[HS95] Zheng-Xu He and O. Schramm, *Hyperbolic and parabolic packings*, Discrete Comput. Geom. **14** (1995), no. 2, 123–149, DOI 10.1007/BF02570699. MR1331923 (96h:52017)

[HR93] Craig D. Hodgson and Igor Rivin, *A characterization of compact convex polyhedra in hyperbolic 3-space*, Invent. Math. **111** (1993), no. 1, 77–111, DOI 10.1007/BF01231281. MR1193599 (93j:52015)

[Koe08] Paul Koebe, *Uber die Uniformisierung beliebiger analytischer Kurven, III*, Nachr. Ges. Wiss. Gott. (1908), 337–358 (German).

[Koe36] Paul Koebe, *Kontaktprobleme der Konformen Abbildung*, Ber. Verh. Sächs. Akad. Wiss. Leipzig **88** (1936), 141–164 (German).

[KMT06] Sadayoshi Kojima, Shigeru Mizushima, and Ser Peow Tan, *Circle packings on surfaces with projective structures: a survey*, Spaces of Kleinian groups, London Math. Soc. Lecture Note Ser., vol. 329, Cambridge Univ. Press, Cambridge, 2006, pp. 337–353. MR2258757 (2008a:52027)

[KK67] G. Kreisel and J.-L. Krivine, *Elements of mathematical logic. Model theory*, Studies in Logic and the Foundations of Mathematics, North-Holland Publishing Co., Amsterdam, 1967. MR0219380 (36 #2463)

[McC96] Gareth McCaughan, *Some results on circle packings*, Ph.D. thesis, University of Cambridge, 1996. At the time of writing, this was available online at `http://www.mccaughan.org.uk/g/personal/maths.html`.

[Moh93] Bojan Mohar, *A polynomial time circle packing algorithm*, Discrete Math. **117** (1993), no. 1-3, 257–263, DOI 10.1016/0012-365X(93)90340-Y. MR1226147 (94h:52038)

[Rat06] John G. Ratcliffe, *Foundations of hyperbolic manifolds*, 2nd ed., Graduate Texts in Mathematics, vol. 149, Springer, New York, 2006. MR2249478 (2007d:57029)

[Riv94] Igor Rivin, *Euclidean structures on simplicial surfaces and hyperbolic volume*, Ann. of Math. (2) **139** (1994), no. 3, 553–580, DOI 10.2307/2118572. MR1283870 (96h:57010)

[Riv96] Igor Rivin, *A characterization of ideal polyhedra in hyperbolic* 3*-space*, Ann. of Math. (2) **143** (1996), no. 1, 51–70, DOI 10.2307/2118652. MR1370757 (96i:52008)

[Riv03] Igor Rivin, *Combinatorial optimization in geometry*, Adv. in Appl. Math. **31** (2003), no. 1, 242–271, DOI 10.1016/S0196-8858(03)00093-9. MR1985831 (2004i:52005)

[RS87] Burt Rodin and Dennis Sullivan, *The convergence of circle packings to the Riemann mapping*, J. Differential Geom. **26** (1987), no. 2, 349–360. MR906396 (90c:30007)

[Roh11] Steffen Rohde, *Oded Schramm: from circle packing to SLE*, Ann. Probab. **39** (2011), no. 5, 1621–1667, DOI 10.1007/978-1-4419-9675-6_1. MR2884870

[Rud87] Walter Rudin, *Real and complex analysis*, 3rd ed., McGraw-Hill Book Co., New York, 1987. MR924157 (88k:00002)

[Sac94] Horst Sachs, *Coin graphs, polyhedra, and conformal mapping*, Discrete Math. **134** (1994), no. 1-3, 133–138, DOI 10.1016/0012-365X(93)E0068-F. Algebraic and topological methods in graph theory (Lake Bled, 1991). MR1303402 (95j:52020)

[Sch91] Oded Schramm, *Rigidity of infinite (circle) packings*, J. Amer. Math. Soc. **4** (1991), no. 1, 127–149, DOI 10.2307/2939257. MR1076089 (91k:52027)

[Sch92] Oded Schramm, *How to cage an egg*, Invent. Math. **107** (1992), no. 3, 543–560, DOI 10.1007/BF01231901. MR1150601 (93c:52009)

[Sch87] E. Schulte, *Analogues of Steinitz's theorem about noninscribable polytopes*, Intuitive geometry (Siófok, 1985), Colloq. Math. Soc. János Bolyai, vol. 48, North-Holland, Amsterdam, 1987, pp. 503–516. MR910731 (89a:52021)

[Ste28] E. Steinitz, *Über isoperimetrische Probleme bei konvexen Polyedern*, J. Reine Angew. Math. **159** (1928), 133–143.

[Ste05] Kenneth Stephenson, *Introduction to circle packing: the theory of discrete analytic functions*, Cambridge University Press, Cambridge, 2005. MR2131318 (2006a:52022)

[Str51] Kurt Strebel, *Über das Kreisnormierungsproblem der konformen Abbildung*, Ann. Acad. Sci. Fennicae. Ser. A. I. Math.-Phys. **1951** (1951), no. 101, 22 (German). MR0051934 (14,549j)

[Swa05] Richard Swan, *Tarski's Principle and the elimination of quantifiers*, 2005, unpublished. At the time of writing, this was available online at `http://www.math.uchicago.edu/~swan/expo/Tarski.pdf`.

[Thu80] William Thurston, *The Geometry and Topology of Three-Manifolds*, Princeton University, 1980, unpublished lecture notes, version 1.1. At the time of writing, these notes were available online at `http://library.msri.org/books/gt3m/`.

[Tut63] W. T. Tutte, *How to draw a graph*, Proc. London Math. Soc. (3) **13** (1963), 743–767. MR0158387 (28 #1610)

[Wil03] G. Brock Williams, *Noncompact surfaces are packable*, J. Anal. Math. **90** (2003), 243–255, DOI 10.1007/BF02786558. MR2001072 (2004h:30055)

[Wil06] G. Brock Williams, *A circle packing measurable Riemann mapping theorem*, Proc. Amer. Math. Soc. **134** (2006), no. 7, 2139–2146 (electronic), DOI 10.1090/S0002-9939-06-08200-1. MR2215785 (2006m:52043)

[Zie95] Günter M. Ziegler, *Lectures on polytopes*, Graduate Texts in Mathematics, vol. 152, Springer-Verlag, New York, 1995. MR1311028 (96a:52011)